

Causal Models for Estimating the Effects of Weight Gain on Mortality
James Robins, MD

Correspondence: James Robins, MD, Mitchell L. and Robin LaFoley Dong
Professor of Epidemiology, Departments of Epidemiology and Biostatistics, Har-
vard School of Public Health, Kresge Building, Room 821, 677 Huntington Av-
enue, Boston, MA 02115; Tel: 617-432-0207; Fax: 617-432-1884; robins@hsph.harvard.edu
Running Head: Causal Models

Abstract: Suppose, contrary to fact, in 1950, we had put the cohort of 18 year old non-smoking American men on a stringent mandatory diet that guaranteed that no one would ever weigh more than their baseline weight established at age 18. How would the counter-factual mortality of these 18 year olds have compared to their actual observed mortality through 2007? We describe in detail how this counterfactual contrast could be estimated from longitudinal epidemiologic data similar to that stored in the electronic medical records of a large HMO by applying g-estimation to a novel structural nested model. Our analytic approach differs from any alternative approach in that in that, in the absence of model misspecification, it can successfully adjust for (i) measured time-varying confounders such as exercise, hypertension and diabetes that are simultaneously intermediate variables on the causal pathway from weight gain to death and determinants of future weight gain, (ii) unmeasured confounding by undiagnosed preclinical disease (i.e reverse causation) that can cause both poor weight gain and premature mortality [provided an upper bound can be specified for the maximum length of time a subject may suffer from a subclinical illness severe enough to affect his weight without the illness becomes clinically manifest], and (iii) the presence of particular identifiable subgroups, such as those suffering from serious renal, liver, pulmonary, and/or cardiac disease, in whom confounding by unmeasured prognostic factors so severe as to render useless any attempt at direct analytic adjustment. However (ii) and (iii) limit the ability to empirically test whether the structural nested model is misspecified. The other two g-methods - parametric g-computation algorithm and inverse probability of treatment weighted (IPTW) estimation of marginal structural models (MSMs) can adjust for potential bias due to (i) but not due to (ii) or (iii).

Key words: BMI, confounders, G-estimation, reverse causation, structural nested failure time model

1 Introduction

Suppose, contrary to fact, in 1950, we had put the cohort of 18 year old non-smoking American men on a stringent mandatory diet that guaranteed that no one would ever weigh more than their baseline weight established at age 18. Specifically, each subject was weighed every day starting on the day before his 18th birthday. Whenever his weight was greater than or equal to this baseline weight, the subject's caloric intake was restricted, without changing his usual

mix of calorie sources and micronutrients, until the time (usually within 1-3 days) that the subject fell to below baseline weight. (I restrict to men solely to avoid the complicating issue of how much weight gain to allow during pregnancy.) Thus, ignoring errors of a pound or two, no subject would ever weigh more than his baseline weight. No instructions or restrictions were given concerning exercise at any time or the amount or nature of what the subject ate during non-calorie restricted periods. How would the counter-factual mortality of these 18 year olds have compared to the actual observed mortality through 2007.

Factually, a substantial fraction of 18 year old American male gains more than 30 lbs from age 18 to 74. Thus if the counterfactual mortality were much less than the observed mortality, then, it would make sense for individuals to maintain their baseline body weight by restricting caloric intake (regardless of whether or not a practical, non-mandatory public health intervention exists that would successfully maintain the baseline weight of most of the (non-smoking) US population.) Here and through we use the phrase "maintain their age x bodyweight" to mean that after age x a subject's weight never exceeds his weight at age x , although it may drop below that weight.

The difference between the counterfactual mortality were no one to exceed their age 18 body weight and the actual observed mortality of the non-smoking US population has been discussed by Willett et al (1) as a useful way to conceptualize the effect of weight on mortality. A major goal of this paper is to show that g-estimation of a novel structural nested model (SNM) can be used to directly estimate this difference from longitudinal observational data. A SNM is a model that takes as input a subject's observed outcome in their observed exposure (here, weight) history, and an unknown parameter and outputs the response that would have been observed if, contrary to fact, the subject to follow the stringent mandatory diet described above. The unknown parameter vector of a SNM is estimated via the g-estimation procedure introduced in Robins et al (12). Previous analytic approaches to the estimation of the effect of weight on mortality do not provide a direct estimate of this difference. In addition, previous approaches have suffered from one or more of the following sources of bias : (i) failure to adequately control for measured confounding due to time-varying exercise, blood lipids, blood pressure, diabetes, and other chronic diseases (once diagnosed) because of concerns that one will thereby be controlling for intermediate variables on the causal pathway from overweight to death, (ii) failure to adequately control for unmeasured confounding due to undiagnosed chronic disease such as cancer (i.e., reverse causation) and (iii) failure to update the weight of a subject whose weight changes after start of follow-up, because of concerns about reverse causation and measurement error.

Bias due to confounding by measured time-varying confounders that are also intermediate variables can be controlled by the use so-called g-methods. G - methods are statistical methods specifically designed to control bias attributable to time-varying confounders affected by previous exposure. In addition to g-estimation of structural nested models, g-methods include the parametric g-formula estimator and inverse probability of treatment weighted (IPTW)

estimators (7,15). As yet g-methods have not been used to estimate the effect of overweight on obesity with the exception of Ref. (2), where the parametric g-formula estimator was used. In this paper we concentrate on g-estimation of SNMs, because as discussed below, of the three g-methods, only g-estimation of SNMs can adjust for unmeasured confounding due to undiagnosed chronic disease.

Finally g-estimation of SNMs allows one to update the weight of a subject whose weight changes after start of follow-up without introducing any bias due to reverse causation. However issues of measurement error are more tricky and will be discussed in the final section of the paper.

Even if maintenance of age 18 weight improves mortality, perhaps a mandatory intervention that allowed weight gain of 0.3/12 pounds per month (i.e., 3 pounds per decade) would produce an even lower mortality. Perhaps the mandatory intervention that would produce the lowest mortality (i.e., the optimal intervention among all "weight-gain" interventions) is one that allows a weight gain of 0.3/12 pounds per month in subjects free of hypertension, diabetes, hyperlipidemia, or clinical CHD, but of only 0.1/12 pounds per two months (i.e., 1 pound per decade) once a subject developed one of these risk factors.

To decide which mandatory intervention is optimal, we require a well-defined numerical measure of overall mortality that can be used to rank interventions. For example, one might use the total years of life (or quality adjusted life) experienced by the cohort from 1950-2007 as a measure. Use of this measure is mathematically equivalent to the use of "years (or quality-adjusted years) of life lived from 1950-2007" as the (subject-specific) utility function in a decision problem whose goal is to maximize expected utility. "Years (or quality-adjusted years) of life lived" measures have a much more natural and useful public health and policy interpretation than the rate ratio, attributable fraction, and attributable risk measures routinely reported in epidemiologic studies.

However even "years of life lived from 1950-2007" is an inadequate utility function when follow-up of the cohort is not to extinction. This function inappropriately assigns the same utility not only to all subjects alive at age 74 on Jan 1, 2008 regardless of their state of health, but also to a subject who dies on Dec. 31, 2007 at 11:59 pm. Clearly among survivors in 2008, the healthier ones (according to some agreed on standard measure of current health) have a greater post-study expected survival (and thus warrant a higher utility) than the less healthy survivors and a much greater expected survival (and thus warrant a much higher utility) than the non-survivors who died in late December 2007. We will not discuss further precisely how to decide on an appropriate utility measure for the survivors, except to remark that such a discussion is necessary. Rather, we will simply assume that, at the end of follow-up, each cohort member has been given a utility measure Y .

Note that the benefit of any of the above counterfactual interventions is an overall effect of the intervention. For example it is conceivable that the mortality benefit of the intervention that maintained baseline weight was wholly due to changes in exercise. Perhaps maintenance of baseline weight makes individuals

feel so much better that they exercise more.

In section 2, we assume we have observational retrospective follow-up data through 2007 on a random sample of the cohort of US males who were non-smokers and 18 in 1950. The data includes detailed medical records, analogous to those currently available on subscribers to a comprehensive HMO. In Sec. 2.2, I discuss three major sources of potential bias that complicate any attempt to estimate the overall effect of the mandatory intervention "maintain baseline weight" on the expected utility of our cohort: (i) measured time-varying confounders such as exercise, hypertension and diabetes that are potentially intermediate variables, (ii) unmeasured confounding by undiagnosed preclinical disease (i.e reverse causation) that can cause both poor weight gain and premature mortality, and (iii) the presence of particular identifiable subgroups, such as those suffering from serious renal, liver, pulmonary, and/or cardiac disease, in whom confounding by unmeasured prognostic factors is so severe as to render useless direct analytic adjustment for confounding. In Section 3, I describe how g-estimation of a correctly specified SNM can appropriately adjust for these potential sources of bias, [provided an upper bound can be specified for the maximum length of time a subject may suffer from a subclinical illness severe enough to affect his weight before the illness becomes clinically manifest]. The SNM required for this adjustment is novel in two ways. First it is a joint SNM, combining a structural nested failure time model (SNFTM) for the counterfactual time to the earlier of death or the diagnosis of a chronic illness and a conditional structural nested mean model (SNMM) for the counterfactual mean of a subject's counterfactual utility given his counterfactual time to death or a diagnosed chronic illness. Second our SNM only models the causal effect of an any increase in BMI between month m and $m + 1$ over a subject's maximum previous BMI. In particular, it does not model and thus is agnostic about the causal effect a) of any decrease in BMI or b) of any increase in BMI between m and $m + 1$ that fails to attain the previous maximum. As a consequence, our SNM is more robust than standard SNMs that also model a) and b), because our model makes fewer assumptions than such alternative models, and thus is less likely to be misspecified. However, the small number of assumptions made by our SNM are sufficient to consistently estimate our parameter of interest $E[Y_0]$.

In Sections 3.2.4-3.2.5, however, I show that (ii) and (iii) limit the ability to empirically test whether the joint structural nested model is misspecified. I also show that, somewhat remarkably, to adjust for bias due to reverse causation one need not assume a deterministic rank-preserving SNM. This is important since a deterministic rank-preserving SNM assumes that the effect of weight-gain on mortality is the same for different subjects, an assumption that is clearly biologically implausible. In Section 4, I consider how to account for censoring by administrative end of follow-up. In Section 5, I consider the estimation of the expected utility under alternative dietary interventions. In Section 6, I discuss the consequences of measurement error in BMI. Proofs and statements of several new theorems are collected in Appendices 1 and 2. Finally, estimation of the optimal "weight-gain" intervention is discussed in the Appendix 3.

2 Estimation of an overall effect

2.1 The Data

I first describe the observational data that is supposed to be available. First, I suppose that a subject's BMI is recorded at the end of each month t , $t = 0, 1, \dots, K$, where *time* t is in months since age 18 and $K+1 = (2007 - 1950) \times 12$ is the duration of follow-up. Let $A^*(t)$ be the difference between BMI at the end of month t and at the end of month $t - 1$. Let $L(t)$ be the vector of covariates values recorded in month t and suppose $L(t)$ precedes $A^*(t)$ temporally. $L(t)$ includes blood pressure, HDL and LDL measures of cholesterol, any diagnoses of and clinical and laboratory characteristics of any chronic disease such as cancer, CAD, diabetes, asthma, COPD, liver, renal disease, etc., level of exercise, measures of mobility and disability, etc. The vector $L(t)$ also includes $BMI(t)$, the BMI just before t rounded to the nearest pound. Thus

$$A^*(t) = BMI(t+1) - BMI(t). \quad (1)$$

$L(t)$ also includes the indicator $I(T > t)$ of vital status at the beginning of month t with T the death time of a subject and, for any proposition B , $I(B)$ is the indicator function that take the value 1 if B is true and zero otherwise. Thus $I(T > t) = 1$ if a subject is alive at t and zero if dead at t . If $I(T > t) = 0$, I include in $L(t)$ the exact day of death and, by convention, assign the value zero to all other components of $L(t)$.

By convention, set $A^*(t)$ and the remaining components of $L(t)$ to zero once a subject has died.

The baseline covariates $L(0)$ include covariate and BMI data on a subject before follow up starts at age 18. Specifically, let $BMI(0)$ denote BMI at (just before) age 18 (i.e., time 0). Our inclusion of BMI just before age 18 as a covariate rather than a treatment reflects the fact that "change" in BMI since 18 is our exposure. In particular, note that $A^*(0)$ is the difference between BMI recorded just before 18 yrs and 1 month and BMI recorded just before 18 years. As is standard in the literature, I have taken change in BMI rather than in change in weight in pounds as the exposure variable. Let $\overline{A^*}(t)$ and $\overline{L}(t)$ be change in BMI and covariate history through time t and $\overline{A^*} = \overline{A^*}(K)$ be a subjects (change in) BMI history through month K and $\overline{L} = \overline{L}(K+1)$ be L history through the end of the study. A subject's utility Y , a measure of quality-adjusted survival, is calculated from $\overline{L} = \overline{L}(K+1)$ since \overline{L} includes the survival time of nonsurvivors, health status measures for survivors at end of follow-up, and time-varying health status factors.

2.2 Potential for Measured and Unmeasured Confounding:

2.2.1 Reverse Causation and Unmeasured Confounding by Subclinical Disease:

In the literature on the effect of BMI on mortality, a controversy has arisen as whether and how to modify standard analytic methods to account for "reverse causation". Reverse causation refers to the well-accepted fact that preclinical (i.e., undiagnosed) chronic disease, such as preclinical cancer, can cause both weight loss (or diminished weight gain) and death. It follows that among subjects with identical BMI history $(\overline{A^*}(t-1), BMI(0))$ and measured covariate history $\overline{L}(t)$ before age t , the subset whose monthly change $A^*(t)$ in BMI is negative *are not comparable* with regard to mortality risk to the subset with positive $A^*(t)$, even if BMI has no causal effect on mortality. That is, reverse causation implies unmeasured confounding by undiagnosed chronic disease. In fact, by an analogous argument, even among the subset with $A^*(t)$ positive, there will be unmeasured confounding, because those with a small gain in BMI are more likely to have preclinical disease than those with a substantial gain.

It follows that one requires an analytic method that can adjust for unmeasured confounding due to the presence of preclinical disease. I will present a method that is appropriate under the additional assumption that we are able to specify an upper bound on the length of time a subject may have a subclinical illness severe enough to affect his weight, before that illness becomes clinically manifest.

2.2.2 Measured Confounders that are also Intermediate Variables

I next turn to the issue of confounding by measured factors, i.e., by components of the covariate vector $L(t)$. For pedagogic purposes, in this subsection, it will be simpler to imagine that the unmeasured confounding due to reverse causation discussed above is not present. Now it is fairly well accepted that obesity causes increased blood pressure (BP), increased low density lipoproteins (LDL), diabetes (Db), and decreased exercise and these four factors may in turn cause increased mortality. Thus these four variables are intermediate variables on the causal pathway from BMI to mortality. In order to prevent underestimation of the overall effect of BMI on mortality due to adjusting for intermediate variables, many analyses of the effect of BMI on mortality have failed to adjust for BP, LDL, Db, or exercise in the analysis. However such a decision can only be justified if these potential intermediate variables do not also confound the BMI-mortality relationship.

A sufficient condition for these intermediate variables to also be confounders is that, among subjects with identical BMI history $(BMI(0), \overline{A^*}(t-1))$ until t , the subset whose monthly change $A^*(t)$ in BMI is negative *are not comparable* with regard to past BP, LDL, Db, and exercise history to the subset with positive $A^*(t)$. Such non-comparability implies that, if data on time-varying BP, LDL, Db, and exercise history are not used in the analysis, there will exist a non-causal

association between an increase of $A^*(t)$ in BMI during month t and subsequent adverse *mortality*, even under the null hypothesis of no overall effect of BMI on mortality. Such non-comparability can occur whenever some or all of these intermediate variables are either a cause of a change in BMI or are correlated with an unmeasured cause of a change in BMI

For example, it is likely that lack of exercise causes weight gain. In that case, if regular exercise causes decreased mortality, then, in an analysis that fails to adjust for exercise history prior to t , the association between an increase of $A^*(t)$ in BMI during month t and subsequent adverse *mortality* will be an overestimate of the true causal effect of $A^*(t)$ on mortality, due to uncontrolled confounding by exercise.

Similarly, suppose that chronic emotional stress and low grade depression not only cause weight gain by inducing overeating as a soothing, self-medicating behavior, but also directly cause elevated BP, elevated LDL, and Db independently via various stress-induced metabolic, immune, and sympathetic nervous system effects. If, as is true in most observational data bases, data on chronic emotional stress and low grade depression are not recorded (i.e., measured), then, even under the null hypothesis of no overall effect of BMI on mortality, the association between an increase $A^*(t)$ in BMI and subsequent adverse *mortality* will tend to be positive, whether or not one adjusts for elevated BP, elevated LDL, and Db in the analysis, due to uncontrolled confounding by chronic emotional stress and low grade depression. However, these variables should be adjusted for in the analysis, because the magnitude of positive overestimation will often be much less if they are adjusted for, because of their correlation with the unmeasured causal confounder - chronic emotional stress and depression.

In contrast with the last paragraph, suppose there is no confounding by chronic emotional stress and low grade depression; rather, in the observational data base, most individuals who developed an elevated BP, elevated LDL, or Db became concerned about their health and instituted a diet that resulted in their gaining less weight than those without these conditions. Then the association found between an increase $A^*(t)$ in BMI and subsequent adverse *mortality* in an analysis that fails to adjust for these variables at t would tend to underestimate the true causal effect of *BMI* on mortality due to negative confounding. I conclude that elevated BP, elevated LDL, and Db could confound the association between increase in BMI and subsequent adverse *mortality* in either a negative or positive direction, depending on which of the mechanisms described in this paragraph and the last predominates.

In summary, time-dependent covariates such as exercise (i.e., physical activity), BP, LDL, or Db that are recorded in $L(t)$ may be both intermediate variables on the causal pathway from BMI to death and confounders of the BMI-death relationship. It follows that one requires an analytic method that can appropriately adjust for the effects of measured time-varying covariates that are simultaneously intermediate variables and time-dependent confounders.

2.2.3 Intractable Unmeasured Confounding in Subgroups

There may be subgroups defined by measured variables in whom confounding by unmeasured factors is intractable. For example, among persons with diagnosed chronic renal, liver, pulmonary or cardiac disease, rapid weight gain can indicate increasing edema (water retention) due to unmeasured disease progression rather than increasing fat stores; as a consequence, among chronic disease patients with identical pasts, comparability would not hold because individuals experiencing rapid weight gain may be at increased risk of death due to unmeasured progression of disease compared to those with lesser weight gain. In such a case unmeasured confounding by disease progression may be intractable.

Using other arguments, various investigators have argued that in both the subgroup of subjects over age 70 and the subgroup with BMI less than 21, subjects gaining weight at different rates are not comparable owing to unmeasured confounding factors, even when data has been collected on many potential confounders. .

Therefore one needs an analytic method that can remain valid even when there exists intractable confounding among subjects with a diagnosed chronic disease, an age of greater than 70, or a BMI below 21. In the next section, I describe an analytic method that satisfies the requirements of this and the two previous subsections.

2.3 A Simplified Description of G-estimation of Structural Nested Models (SNMs):

In this subsection I give a nontechnical, conceptual description of how, even in the presence of the measured and unmeasured confounding described in Section 2.2, g-estimation of structural nested models can be used to estimate the expected utility had, contrary to fact, all non-smoking 18 year old American men in 1950 been put on a stringent mandatory diet that guaranteed that no one would ever weigh more than their weight at age 18. In order to avoid technical digressions and thereby keep the description centered on important conceptual issues, this nontechnical description is neither complete nor fully accurate. Section 3 onwards provides a complete and accurate description. This completeness and accuracy unfortunately place greater technical demands on the reader.

A locally rank preserving SNM for Y is a rule that takes as input a subject's observed utility Y , their observed BMI and covariate history through the end of the study, and an unknown parameter β^* and outputs the utility Y_0 that would have been observed if, possibly contrary to fact, the subject had followed the dietary intervention of the first paragraph of the Introduction. If the rule is correct and we knew the value of β^* , then we could calculate Y_0 for each study subject. The average of these Y_0 in the cohort of all non-smoking 18 year old American men in 1950 is our quantity of interest: the expected (i.e. average) utility had one implemented a dietary intervention that guaranteed that no one would ever weigh more than they did at age 18. However we do not know the

value of β^* . Thus the challenge is to estimate β^* from the data. When, as in section 2.2.2, all confounding is due to measured variables, Robins (12) proposed a method of estimation called g-estimation that is described next.

If the only confounding is due to measured factors, then among subjects with the same BMI and covariate history prior to time t with nonnegative $A(t)$, the increase $A(t)$ in BMI between t and $t+1$ will be conditionally uncorrelated with Y_0 . Thus to estimate β^* , we simply try many different guesses β . If a particular guess β were the true β^* , then the output of the rule would be uncorrelated with $A(t)$. Thus I choose as our estimate $\hat{\beta}$ of β^* , the guess β which results in an output that has smallest conditional correlation with $A(t)$ when we combine the information across all months t from 0 to end of follow-up at $K+1$.

When as in Section 2.2.3, there are certain identifiable subgroups in whom confounding is intractable, bias can result because the output of the rule will be conditionally correlated with $A(t)$ even when $\beta = \beta^*$. To eliminate this bias, it suffices to search for lack of correlation with $A(t)$ only among the subset of subjects who are not members of these intractable confounded subgroups at time t . That is, we simply restrict our g-estimation procedure at a given time t to subjects who are not currently members of these subgroups.

When there is unmeasured confounding by subclinical disease such as in Section 2.2.1, I must modify our g-estimation procedure. Suppose one can specify an upper bound, say 6 years, on the length of time a subject may have a subclinical illness severe enough to affect weight gain, before that illness becomes clinically manifest. Then one can still validly estimate β^* if one restricts the g-estimation procedure at a given time t to those subjects who would have remained alive and free of a diagnosed chronic (i.e., of clinical) disease for the six years following t had, possibly contrary to fact, they followed a diet that prevented any further weight gain over those 6 years; by our assumption of a 6-year upper bound, such subjects did not have their weight gain affected by an undiagnosed chronic disease. [It does not suffice to restrict to subjects who actually remained alive and free of clinical disease for the six years following t , because if BMI change $A(t)$ at t causally effects the onset of clinical disease and/or survival in the following six years, the variable 'survival without clinical disease for six years after t ' is a response affected by the exposure $A(t)$ and thus cannot be adjusted for without introducing selection bias as explained in Hernan et al. (13). Thus to validly estimate β^* using g-estimation, one must be able to determine those "subjects who would have remained alive and free of clinical disease for the six years following t had, possibly contrary to fact, they followed a diet that prevented any further weight gain over those 6 years."

One can do so by specifying a second SNM, called a locally rank preserving structural nested failure time model (SNFTM), for the effect of change in BMI on the time X to the diagnosis of chronic disease or death (whichever comes first). A locally rank preserving SNFTM is a rule that takes as input a subject's observed time X to (the earlier of) death or a diagnosed chronic disease, their observed BMI and covariate history through the end of the study, and an unknown parameter ψ^* , and a time t and outputs the time X_t that would have been observed if, possibly contrary to fact, the subject had followed a di-

etary intervention in which no further weight was gained after time t . If ψ^* were known or well-estimated, we could compute X_t for each subject, determine which subjects' X_t failed to exceed t by more than 6 years, and exclude such subjects from the g-estimation procedure used to estimate β^* .

Thus it only remain to estimate the parameter ψ^* of our locally rank preserving SNFTM in the prescence of unmeasured confounding by subclinical disease. Now among subjects with the same BMI and covariate history prior to time t with nonnegative $A(t)$ who are not members of an identifiable subgroup with intractable confounding, the change $A(t)$ in BMI between t and $t + 1$ will be uncorrelated with X_t if we restrict to subjects with X_t exceeding t by more than 6 years. Thus to estimate ψ^* , I simply try many different guesses ψ . If a particular guess ψ were the true ψ^* , the output of the SNFTM rule would be uncorrelated with $A(t)$ when I restrict the g-estimation procedure to subjects whose output exceeds t by more than 6 years. Thus I choose as the estimate $\hat{\psi}$ of ψ^* , the guess ψ which results in an output that, under this restricted g-estimation procedure, has the smallest conditional correlation with $A(t)$ when I combine the information across all times t from 0 to *end* of follow-up at $K + 1$.

Before proceeding to the more technical part of the paper, I provide a brief non-technical discussion of several important but subtle points about SNMs. First, locally rank preserving SNMs assume that the effect of a given increase in BMI on the utility Y and on X is the same for any two subjects with the same past measured covariate history. This assumption is biologically implausible since unmeasured genetic and enviornmental factors will clearly modify the magnitude of the effect of weight gain on the responses Y and X . Fortunately, we prove in Section 3 that our g-estimator of the mean of Y_0 remains valid even if we allow the magnitude of the effect of weight gain on Y and X to be modified in an arbitrary manner by unmeasured genetic and enviornmental factors.

The description of g-estimation of the parameters ψ^* of our SNFTM model for X assumed that the time X to death or diagnosed chronic disease was available for every study subject. However, by end of follow-up, a number of study subjects will remain alive and free of chronic disease. Such subjects are said to be censored. In Section 4, I show how our g-estimation procedures can be modified to appropriately account for these censored observations.

The estimate of the mean of Y_0 will be biased if either the SNMM for Y or the SNFTM for X are misspecified. I discuss below how to construct tests for misspecification. However, I also show that the power of such tests to detect model misspecification can be quite limited in the prescence of reverse causation by subclinical disease and intractable confounding in identifiable subgroups. In Section 3.3, I offer some suggestions on how the impact of this limited power on the quality of one's inferences can be lessened if one is willing to change the parameter that is being estimated.

3 Estimation of the effect of the "maintain baseline weight intervention"

In this section I describe how we can use G-estimation of structural nested models (SNMs) to estimate the the expected utility had one put all non-smoking 18 year old American men on a stringent mandatory diet that guaranteed that no one would ever weigh more than their baseline weight established at age 18. For pedagogic reasons I first consider the simpler setting in which there is no unmeasured confounding by preclinical disease.

3.1 Case 1: No unmeasured confounding by preclinical disease.

3.1.1 A Locally Rank Preserving SNM.

An SNM is a model for counterfactual variables Y_m that denote a subject's utility measured at end of follow-up under the following counterfactual dietary intervention:

Time m Dietary Intervention: The subject follows his observed diet up to month m following his 18th birthday and, from month m onwards, the subject is weighed every day: (i) whenever his weight is greater than or equal to his maximum monthly BMI up to m [i.e., $BMI_{\max}(m) \equiv \max\{BMI(0), \dots, BMI(m)\}$], the subject's caloric intake is restricted until the subject's BMI falls to below $BMI_{\max}(m)$; (ii) whenever his weight is less than $BMI_{\max}(m)$, the subject is allowed to eat as he pleases without any intervention.

A subject's responses had, possibly contrary to fact, he been made to follow a time m dietary intervention are referred to as counterfactual responses. We assume that Y_m is well-define in the sense that its value is insensitive to the unspecified details of exactly how the subject's calories are to be restricted in (i). We also assume a subject's counterfactual responses are observed only for those m for which a subject's actual BMI history was consistent with his having followed the time m dietary intervention. For other values of m , the time m -specific counterfactuals remain unobserved.

The time 0 dietary intervention is the dietary intervention in the first paragraph of the Introduction. The counterfactual Y_0 is the utility corresponding to this regime. Thus the expected value $E[Y_0]$ of Y_0 is our parameter of interest: the expected utility had we placed in 1950 all non-smoking 18 year old American men on a diet that guaranteed that no one would ever weigh more than they did at age 18.

Note that $Y_{K+1}=Y$: if one were to follow his actual observed diet up to the time $K+1$ at which the study ends, then no dietary intervention would have occurred. Hence the counterfactual Y_{K+1} must be the observed (i.e., actual) Y .

By definition, a subject's observed data through k (but before $k+1$) is inconsistent or incompatible with following the "time m dietary intervention" if and only if $BMI(k+1) > BMI_{\max}(m)$ for some $k > m$.

Let us define $A(t)$ to be the difference between a subject's observed BMI, $BMI(t+1)$, just prior to month $t+1$ and his maximum value $BMI_{\max}(t)$ of BMI prior to month t , whenever that difference is nonnegative. When the difference is negative, we simply set $A(t)$ to be zero. Formally then

$$A(t) = BMI(t+1) - BMI_{\max}(t) \text{ if } BMI(t+1) \geq BMI_{\max}(t) \quad (2)$$

$$A(t) = 0 \text{ if } BMI(t+1) < BMI_{\max}(t). \quad (3)$$

$A(t)$ is nonnegative. It follows that it is only when the individual's observed data is incompatible with the "time m dietary intervention" through time m is $A(m) \neq 0$. If an individual's observed data is consistent with his having followed the "time m dietary intervention", it is consistent with his having followed the "time t dietary intervention" for $t > m$.

Note that $Y_{m+1} - Y_m = 0$ whenever $A(m) = 0$. If $A(m) \neq 0$, $Y_{m+1} - Y_m$ is the difference between (i) a subject's utility when he has his observed $\overline{BMI}(m+1)$ history and thereafter, possibly contrary to fact, the subject follows the dietary intervention that guarantees his $BMI(k)$ for $k > m+1$ never again exceeds $BMI(m+1)$ and (ii) his utility when he has his observed $\overline{BMI}(m)$ history and thereafter, possibly contrary to fact, the subject follows the dietary intervention that guarantees his BMI at $m+1$ equals his observed $BMI_{\max}(m)$ [rather than his observed BMI at $m+1$] and that his $BMI(k)$ for $k > m+1$ never again exceeds $BMI_{\max}(m)$. As a kind of shorthand for the previous sentence, whenever $A(m) \neq 0$, we will refer to $Y_{m+1} - Y_m$ as the causal effect of final blip of exposure of magnitude $A(m)$ on the subject's utility.

An additive locally rank preserving SNM is a deterministic model for the magnitude of the effect of a treatment $A(m)$ on $Y_{m+1} - Y_m$. Mathematically an additive locally rank preserving SNM assumes that for each time $m = 0, \dots, K$,

$$Y_{m+1} - Y_m = \gamma_m [A(m), \overline{A}(m-1), \overline{L}(m), \beta^*] \quad (4)$$

where (i) β^* is the unknown true parameter vector, and (ii) $\gamma_m [A(m), \overline{A}(m-1), \overline{L}(m), \beta]$ is a known function [such as $\{\beta_0 + \beta_1 m + \beta_2^T L(m)\} A(m)$] satisfying the restrictions $\gamma_m [A(m), \overline{A}(m-1), \overline{L}(m), \beta] = 0$ if $A(m) = 0$ or $\beta = 0$. Here θ_2 is a column vector of length equal to that of the vector $L(m)$. Furthermore, T as a superscript denotes the transpose of a matrix or vector. The first restriction must logically hold because, by definition, if $A(m) = 0$, $Y_{m+1} = Y_m$. We now show that the second restriction guarantees that $\beta^* = 0$ encodes the sharp null hypothesis that "following a diet that prevents one's BMI from ever exceeding the baseline BMI" has no effect on any subject's utility.

Recalling that $Y_{K+1} = Y$, the model (4) is seen to be equivalent to the model

$$Y_m = Y - \sum_m^K \gamma_m [A(m), \overline{A}(m-1), \overline{L}(m), \beta^*] \quad (5)$$

for $m=0,1,\dots,K$. To help understand equation (5) consider first the special case $m = K$. Then equation (5) says that to calculate Y_K from Y , we remove

the causal effect $\gamma_K [A(K), \bar{A}(K-1), \bar{L}(K), \beta^*]$ of exposure $A(K)$ at the last time K . Next consider the special case $m = 0$. Then equation (5) says that to calculate Y_0 , one successively removes the effect of exposure at times $K, K-1, \dots, 0$. It follows from the restriction $\gamma_m [A(m), \bar{A}(m-1), \bar{L}(m), 0] = 0$ that $\beta^* = 0$ implies $\gamma_m [A(m), \bar{A}(m-1), \bar{L}(m), \beta^*] = 0$ for each m . Thus $\beta^* = 0$ encodes the sharp null hypothesis that $Y_0 = Y_m = Y$ for all subjects and all m . In other words, one's utility at the end of the study will be the same regardless of whether or not one follows any "time m dietary intervention".

Since, by Eq. (5), a locally rank preserving SNM directly maps an individual's observed utility Y to the utility an individual would have under the "time m dietary intervention", it is a model for individual causal effects.

Possible choices of $\gamma_m [a(m), \bar{a}(m-1), \bar{l}(m), \beta]$ include (i) $\beta a(m)$,

(ii) $(\beta_0 + \beta_1 m) a(m)$, (iii) $\{\beta_0 + \beta_1 m + \beta_2^T l(m)\} a(m)$. In model (i), the effect of a change of $A(m)$ in BMI is the same for all m . Under model (ii), the effect varies linearly with time m . Under model (iii), the causal effect of $A(m)$ is modified by the most recent covariate history.

In the following we assume the observed data O on each subject is $O = (Y, \bar{L}, \bar{A}) \equiv (Y, \bar{L}(K+1), \bar{A}(K))$. That is O consists of a subject's utility Y and his covariate and treatment histories through the end of the study. The inclusion of $\bar{A}(K)$ is actually redundant, since the A -history $\bar{A}(K)$ is determined by $\bar{BMI}(K+1)$, and $\bar{BMI}(K+1)$ is a component of $\bar{L}(K+1)$. Thus we could write the observed data as simply $(Y, \bar{L}(K+1))$. However because we wish to use results on g-estimation of SNMs that were derived in previous papers in which $\bar{A}(K)$ was not determined by $\bar{L}(K+1)$, we will continue to write $O = (Y, \bar{L}, \bar{A})$ and accept some redundancy in the notation. Let

$$Y_m(\beta) = Y - \sum_{j=m}^K \gamma_m [A(j), \bar{A}(j-1), \bar{L}(j), \beta] \quad (6)$$

so, under our model, $Y_m = Y_m(\beta^*)$. Note that, for each β , $Y_m(\beta)$ can be computed from the observed data (Y, \bar{L}, \bar{A}) . Suppose we had a consistent estimate $\hat{\beta}$ of β^* . Then $Y_0(\hat{\beta})$ would be a consistent estimate of $Y_0 = Y_0(\beta^*)$. Thus the average $\sum_{i=1}^n Y_{0i}(\hat{\beta})/n$ over the n study subjects would be a consistent estimate of the parameter of interest $E[Y_0]$. Further $\sum_{i=1}^n Y_{0i}(\hat{\beta})/n - \sum_{i=1}^n Y_i/n$ would be a consistent estimate of the difference $E[Y_0] - E[Y]$ between the expected utility $E[Y_0]$ under a dietary intervention guaranteeing BMI never exceeds the baseline BMI and the expected utility $E[Y]$ in the absence of any dietary intervention. Below we show how one can obtain a consistent estimate $\hat{\beta}$ by g-estimation if a certain comparability assumption holds.

The Innovative Aspect of our SNM: The most important and innovative aspect of our model is that it models the causal effect on the utility of an increase in BMI of $A(m)$ over a subject's maximum past BMI, $BMI_{\max}(m)$. It does not model and thus is agnostic about the causal effect a) of any decrease in BMI or b) of any increase in BMI between m and $m+1$ that fails to result

in one's BMI exceeding $BMI_{\max}(m)$. Our model (4) is thus more robust than alternative models that would also model a) or b). However, the small number of assumptions made by our model are sufficient for our purposes; if we can consistently estimate the parameter β^* , we can consistently estimate our parameter of interest $E[Y_0]$.

Thus it only remains to estimate β^* . In this section will do so under the following assumption, which will be weakened in later sections. Define the indicator variable $\Xi(m)$ taking values in the two-element set $\{0, 1\}$ by

$$\Xi(m) = 1 \Leftrightarrow BMI(m+1) \geq BMI_{\max}(m) \quad (7)$$

That is $\Xi(m)$ takes the value 1 if a subject's BMI just before $m+1$ is at least as great as his maximum BMI up to time m . Otherwise $\Xi(m)$ takes the value 0.

Comparability Assumption (CO): Among subjects with the same $\bar{A}(m-1)$ history and covariate history $\bar{L}(m)$ (which includes BMI history $\bar{BMI}(m)$) and with $\Xi(m) = 1$, $A(m)$ is statistically independent of the counterfactual Y_m . Formally, conditional on $(\bar{A}(m-1), \bar{L}(m), \Xi(m) = 1)$, $A(m)$ is independent of Y_m . [Since past A -history $\bar{A}(m-1)$ is determined by $\bar{BMI}(m)$, $\bar{A}(m-1)$ in the conditioning event $(\bar{A}(m-1), \bar{L}(m))$ is redundant; nonetheless we shall retain the $\bar{A}(m-1)$.]

A comparability assumption such as CO is often referred to as an assumption of no confounding by unmeasured factors or as an assumption of sequential randomization.

Remark: To understand why we conditioned on $\Xi(m) = 1$ in the CO assumption, imagine we had instead assumed that $A(m)$ is independent of Y_m conditional on $(\bar{A}(m-1), \bar{L}(m))$. That would have implied that among the subset of subjects with a given $(\bar{A}(m-1), \bar{L}(m))$, the subgroup with $A(m) \neq 0$ would have the same distribution of the utility Y_m under the time m - dietary intervention as the subgroup with $A(m) = 0$. But, under the time m - intervention, all subjects in the $A(m) \neq 0$ subgroup would have $BMI(m+1)$ equal to their common $BMI_{\max}(m) \in \bar{L}(m)$, while many subjects with $A(m) = 0$ (specifically, those with $\Xi(m) = 0$) would have $BMI(m+1) < BMI_{\max}(m)$. Thus, the $A(m) = 0$ subgroup will have lower BMI at $m+1$ than the $A(m) \neq 0$ subgroup under the time m - intervention. Suppose the null hypothesis of no biological BMI effect is false. Then, for an individual with $A(m) = 0$, their utility Y_m should depend on their BMI at $m+1$. As such, it is extremely unlikely that the $A(m) = 0$ and $A(m) \neq 0$ subgroups would be comparable. In contrast, if, as in assumption CO, we restrict the $A(m) = 0$ subgroup to a subset of the subjects with $\Xi(m) = 1$, then given $\bar{L}(m)$, this restricted $A(m) = 0$ subgroup, like the $A(m) \neq 0$ subgroup, will have $BMI(m+1)$ equal to the common $BMI_{\max}(m)$ under the intervention, so the assumption of noncomparability is plausible.

It is interesting to note that if we had used the coding convention that vector $L(k)$ includes $\Xi(k)$ as a component, we could then have stated our comparability assumption as $A(m)$ is independent of Y_m conditional on $(\bar{A}(m-1), \bar{L}(m))$ because, under this coding, $pr(A(k) = 0 | \bar{A}(k-1) = \bar{0}(k-1), \bar{L}(k))$ is one

whenever $\Xi(k)$ takes the value zero. However, we will not use this coding convention.

Under the CO assumption, we can obtain a consistent estimator of β^* by g-estimation as follows. We specify a linear regression model

$E[A(m) | \bar{L}(m), \bar{A}(m-1), \Xi(m) = 1] = \alpha^T W(m)$ for $m = 0, \dots, K$. Here $W(m) = w_m [\bar{L}(m), \bar{A}(m-1)]$ is a vector of covariates calculated from a subject's past data, α^T is a row vector of unknown parameters, and each person-month is treated as an independent observation, so each person contributes up to $K + 1$ observations. However, person months for which $\Xi(m) \neq 1$ are excluded from the regression. Examples of $W(m) = w_m [\bar{L}(m), \bar{A}(m-1)]$ would be the transpose of the row vector $(m, L^T(m), L^T(m-1))$. Let $\hat{\alpha}$ be the OLS estimator of α computed using a standard statistical package.

For the moment assume β is one dimensional. Let β_{low} and β_{up} be much smaller and larger, respectively, than any substantively plausible value of β^* .

Then, separately, for each β on a grid from β_{low} to β_{up} , say $\beta_{low}, \beta_{low} + 0.1, \beta_{low} + 0.2, \dots, \beta_{up}$, perform the score test of the hypothesis $\theta = 0$ in the extended linear model

$$E[A(m) | \bar{L}(m), \bar{A}(m-1), Y_m(\beta), \Xi(m) = 1] = \alpha^T W(m) + \theta Y_m(\beta) \quad (8)$$

that adds the covariate $Y_m(\beta)$ at each time m to the above (pooled over persons and time) linear model. A 95% confidence interval for β^* is the set of β for which an $\alpha = 0.05$ two-sided score test of the hypothesis $\theta = 0$ does not reject. The g-estimate $\hat{\beta}$ of β^* is the value of β for which the score test takes the value zero (i.e., the p-value is one).

The validity of g-estimation is proved as follows. By our comparability assumption $Y_m(\beta^*)$ and $A(m)$ are conditionally independent given

$(\bar{L}(m), \bar{A}(m-1), \Xi(m) = 1)$. That is, $Y_m(\beta^*)$ is not a predictor of $A(m)$ given $(\bar{L}(m), \bar{A}(m-1), \Xi(m) = 1)$, which implies that the coefficient θ of $Y_m(\beta)$ must be zero in the extended model when $\beta = \beta^*$, provided the model

$$E[A(m) | \bar{L}(m), \bar{A}(m-1), \Xi(m) = 1] = \alpha^T W(m) \text{ is correctly specified.}$$

Now, we do not know the true value of β . Therefore, any value β for which the data are consistent with the parameter θ of the term $\theta Y_m(\beta)$ being zero might be the true β^* , and thus belongs in our confidence interval. If consistency with the data is defined at the 0.05 level, then our confidence interval will have coverage of 95%. Furthermore, the g-estimate $\hat{\beta}$ of β^* is that β for which adding the term $\theta Y_m(\beta)$ does not help to predict $A(m)$ whatsoever, which is the β for which the score test of $\theta = 0$ is precisely zero. The g-estimate $\hat{\beta}$ is also the value of β for which the OLS estimator of θ is precisely zero.

It may appear peculiar that a function $Y_m(\beta)$ of the response Y measured at end of follow-up is being used to predict $A(m)$ at earlier times. However, this peculiarity evaporates when one recalls that, for each β on our grid, we are testing the null hypothesis that $\beta = \beta^*$, and, under this null, $Y_m(\beta)$ is the counterfactual Y_m , which we can view as already existing at time m (although we cannot observe its value until time $K + 1$ and then only if $A(t)$ in the observed data is zero from m onwards).

Suppose next that the parameter β is a vector. To be concrete suppose we consider the model with

$\gamma_m [a(m), \bar{a}(m-1), \bar{l}(m), \beta] = a(m) \{\beta_0 + \beta_1 m + \beta_3^T l(m)\}$ so β is of dimension $\dim(l(m)) + 2$ where $\dim(l(m))$ is the dimension of $l(m)$ which, for concreteness, we take to be 3. Hence β is 5 dimensional. Then we would use a 5 dimensional grid, one dimension for each component of β . So if we had 20 grid points for each component, we would have 20^5 different values of β on our 5 dimensional grid. Now to estimate 5 parameters one requires 5 additional covariates. Specifically, let $Q_m(\beta) = q_m [\bar{L}(m), \bar{A}(m-1), Y_m(\beta)]$ be a 5 dimensional vector of functions of $(\bar{L}(m), \bar{A}(m-1), Y_m(\beta))$, such as $Q_m(\beta) = [1, m, L^T(m)] (Y_m(\beta))^2$. We use the extended model

$$E[A(m) | \bar{L}(m), \bar{A}(m-1), Y_m(\beta), \Xi(m) = 1] = \alpha^T W(m) + \theta^T Q_m(\beta).$$

Our g-estimate $\hat{\beta}$ is the β for which the 5 degree of freedom score test that all 5 components of θ equal zero is precisely zero. The particular choice of the functions q_m does not affect the consistency of the point estimate, but it affects the width of the confidence interval.

When $\gamma_m [a(m), \bar{a}(m-1), \bar{l}(m), \beta] = a(m) \beta^T R_m$ is linear in β with $R_m = r_m(\bar{L}(m), \bar{A}(m-1))$ being a vector of known functions and we choose $Q_m(\beta) = Q_m^* Y_m(\beta)$ linear in $Y_m(\beta)$, then, given the OLS estimator $\hat{\alpha}^T$ of α^T in the model $E[A(m) | \bar{L}(m), \bar{A}(m-1)] = \alpha^T W(m)$, there is an explicit closed form expression for $\hat{\beta}$ given by

$$\hat{\beta} = \left\{ \sum_{i=1, m=0}^{i=n, m=K} \Xi_i(m) A_i(m) G_{im}(\hat{\alpha}) Q_{im}^* S_{im}^T \right\}^{-1} \left\{ \sum_{i=1, m=0}^{i=n, m=K} \Xi_i(m) Y_i G_{im}(\hat{\alpha}) Q_{im}^* \right\} \quad (9)$$

with $G_{im}(\hat{\alpha}) = [A_i(m) - \hat{\alpha}^T W_i(m)]$, $S_{im} = \sum_{i=1, j=m}^{i=n, j=K} R_{ij}$.

Identification : Suppose that two different values of β , say $\hat{\beta}$ and $\hat{\hat{\beta}}$, both make the 5 degree of freedom score test precisely zero and yet the two CI for β^* centered at $\hat{\beta}$ and $\hat{\hat{\beta}}$ do not overlap. How should we choose between the estimates? [In such a case, the matrix whose inverse is required in (9) will not be invertible and so (9) will fail.] Since we can use any 5 vector $Q_m(\beta) = q_m [\bar{L}(m), \bar{A}(m-1), Y_m(\beta)]$ in our procedure, one simple approach is to try other choices of $Q_m(\beta)$ until we find a $Q_m(\beta)$ for which our CI for β^* includes only one of the $\hat{\beta}$ and $\hat{\hat{\beta}}$, declare the one included to be our point estimate of β^* and ignore the excluded one. Will this approach always succeed? In general this approach should succeed in rather quickly excluding all but one of the values of β that originally made the score test zero, provided that the model $\gamma_m [a(m), \bar{a}(m-1), \bar{l}(m), \beta]$ is correct, except when β^* is not identified. By definition β^* is not identified if there is a β^{**} different from the true parameter β^* such that, with an infinite sample size, β^{**} , like β^* , makes the 5 degree of freedom score test precisely zero for all choices of $Q_m(\beta)$. In our model, it

follows from Robins (3, 4) that under the positivity assumption that

$$\Pr [A(m) = 0 | \bar{A}(m-1), \bar{L}(m), \Xi(m) = 1] \neq 0 \quad (10)$$

for all subjects and $m = 0, \dots, K$, β^* is identified. In our context the positivity assumption is a very weak assumption, that is almost certainly true. Hence, for the remainder of the paper, we will silently assume that it holds.

Remark: We considered a linear regression model for

$E[A(m) | \bar{L}(m), \bar{A}(m-1), Y_m(\beta), \Xi(m) = 1]$ in the above for expositional simplicity. In practice since $A(m) \geq 0$, we might use a log linear model that specifies

$E[A(m) | \bar{L}(m), \bar{A}(m-1), Y_m(\beta), \Xi(m) = 1] = \exp\{\alpha^T W(m) + \theta^T Q_m(\beta)\}$ and fit by non-linear least squares. In that case, in the last display, $G_{im}(\hat{\alpha}) = [A_i(m) - \exp\{\hat{\alpha}^T W_i(m)\}]$. Alternatively we could replace the response variable $A(m)$ in the linear regression by $\ln(A_m + 0.1)$ where the 0.1 is added to insure the logarithm remains finite even when $A_m = 0$. In that case, $G_{im}(\hat{\alpha}) = [\ln\{A_i(m) + 0.1\} - \hat{\alpha}^T W_i(m)]$

3.1.2 An additive structural nested mean model (SNMM)

An additive locally rank preserving SNM (4) implies that if two subjects have the same observed data $O = (Y, \bar{L}, \bar{A})$ they will have the same value of Y_0 under the "time 0 dietary intervention" of the introduction. That is the model implies that for these subjects, the effect of the "time 0 dietary intervention" will be identical. This assumption is clearly biologically implausible in view of between-subject heterogeneity in unmeasured genetic and environmental factors. To overcome this limitation, we consider an additive structural nested mean model (SNMM)

$$E[Y_{m+1} - Y_m | \bar{A}(m), \bar{L}(m)] = \gamma_m [A(m), \bar{A}(m-1), \bar{L}(m), \beta^*] \quad (11)$$

that models the conditional mean of $Y_{m+1} - Y_m$ given $(\bar{A}(m), \bar{L}(m))$ rather than the individual differences $Y_{m+1} - Y_m$, and thus does not impose local rank preservation. In particular Y_m no longer is equal to $Y_m(\beta^*)$. However, Robins (4, 5) proved the additive SNMM implies (and, in fact is equivalent to) the assumption that Y_m and $Y_m(\beta^*)$ have the same mean given $\bar{A}(m), \bar{L}(m), \Xi(m) = 1$. That is

$$E[Y_m | \bar{A}(m), \bar{L}(m)] = E[Y_m(\beta^*) | \bar{A}(m), \bar{L}(m)] \quad (12)$$

for each m . Further he proved that, under the CO assumption, g-estimation of β^* retains all the properties described above, even in the absence of local rank preservation, except now the function $Q_m(\beta)$ must be chosen linear in $Y(\beta)$, i.e., $Q_m(\beta) = Q_m^* Y_m(\beta)$ as above.

As a consequence, the definition of non-identifiability must be modified as follows: the parameter β^* of an SNMM β^* is not identified if there is a β^{**} different from the true parameter β^* such that, with an infinite sample size, β^{**} , like β^* , makes the 5 degree of freedom score test precisely zero for all choices of $Q_m(\beta)$ that are linear in $Y_m(\beta)$. A fuller discussion of rank preserving versus non-rankpreserving models occurs in Section 3.2.4.

Alternative Approaches: Under the CO assumption, it is shown in ref [12] that the $E[Y_m]$ are nonparametrically identified for $m = 0, \dots, K$ from data $O = (Y, \bar{L}, \bar{A})$ by the IPTW formula

$$E[YI\{\underline{A}(m) = \underline{0}(m)\} \mathbb{W}(m)], \quad (13)$$

where $\underline{A}(m) = (A(m), \dots, A(K))$ and the IPTW weight

$$\mathbb{W}(m) = 1 / \prod_{k=m}^K \{pr(A(k) = 0 | \bar{A}(k-1), \bar{L}(k))\}^{\Xi(k)}$$

is the inverse of the conditional probability that a subject had his observed treatment $\underline{A}(m) = \underline{0}(m)$. That is, $E[Y_m]$ is the weighted mean of the observed utility Y among subjects who observed data was consistent with following the time m dietary intervention with weights given by the inverse of the conditional probability of having data consistent with following the intervention. Thus one could, in principle, consider estimating $E[Y_0]$ nonparametrically by the weighted average of Y among subjects whose weight never exceeded their baseline weight at age 18 with weights proportional to an estimate $\widehat{\mathbb{W}}(0)$ of $\mathbb{W}(0)$.

That is, by $\left[\sum_{\{i; \underline{A}_i(0) = \underline{0}(0)\}} \widehat{\mathbb{W}}_i(0) Y_i \right] / \left[\sum_{\{i; \underline{A}_i(0) = \underline{0}(0)\}} \widehat{\mathbb{W}}_i(0) \right]$. The problem with

this approach is that only the utility Y of the rare person whose weight never exceeds his age 18 weight contributes to the analysis. In contrast by specifying a SNMM, data on the utility Y of every subject contributes to the estimate of $E[Y_0]$. The price paid for the greater efficiency of a SNMM is the possibility of bias if the SNMM (11) is misspecified.

However, under the CO, $E[Y_m | \bar{A}(m), \bar{L}(m)] = E[Y_m | \bar{A}(m-1), \bar{L}(m)]$ and is nonparametrically identified by the formula

$E[YI\{\underline{A}(m) = \underline{0}(m)\} / \mathbb{W}(m) | \bar{L}(m), \bar{A}(m-1)]$. Thus $E[Y_{m+1} - Y_m | \bar{A}(m), \bar{L}(m)]$ is nonparametrically identified for $m=0, \dots, K$. Hence, given a sufficiently large sample size, one could in principle construct misspecification tests of the model (11) that have power against all alternatives when the model is incorrect. In practice, the available sample size may greatly limit the power to detect model misspecification.

IPTW estimation of marginal structural models and the parametric g-formula are alternative approaches to model-based estimation of $E[Y_0]$ that also use data on every subject's utility Y . See Appendix 2 for further discussion of the latter approach.

Remark: The reader familiar with IPTW expects $\mathbb{W}(m)$ to be defined as $\mathbb{W}(m) = 1 / \prod_{k=m}^K \{pr(A(k) = 0 | \bar{A}(k-1) = \bar{0}(k-1), \bar{L}(k))\}$ rather than as

$1 / \prod_{k=m}^K \{pr(A(k) = 0 | \bar{A}(k-1) = \bar{0}(k-1), \bar{L}(k))\}^{\Xi(k)}$. In fact, the two expressions would have been equal had we used the coding convention that the vector $L(k)$ includes $\Xi(k)$ as a component because, under this coding, $pr(A(k) = 0 | \bar{A}(k-1) = \bar{0}(k-1), \bar{L}(k))$

is one whenever $\Xi(k)$ takes the value zero. However, we do not use this convention.

3.2 Case 2: Unmeasured confounding by preclinical disease.

In this section we no longer assume $A(m)$ is statistically independent of Y_m given $\bar{A}(m-1), \bar{L}(m), \Xi(m) = 1$. To describe our new comparability assumption we need to introduce some further notation. Let $X = \min(T, \mathcal{D})$ be the minimum of the time T to death and the time \mathcal{D} to the diagnosis of a chronic disease, such as cancer, severe emphysema, liver or renal disease, or any other chronic condition that would be severe enough to affect weight gain. At each time m , the indicator $I(X \leq m)$ is a component of $L(m)$. Further if $I(X \leq m) = 1$, the exact time X is observed and included in $L(m)$. Thus X is observed if X is less than $K+1$. However X is censored (i.e. not observed) on subjects whose X exceeds $K+1$, the end of follow up time. For the present we shall avoid the additional complications that arise from censoring by assuming that X is less than $K+1$ for all subjects, so that the data $O = (Y, X, \bar{L}, \bar{A})$ is observed on each subject. In Section 4, we relax this assumption and allow for censoring.

Let $X_m = \min(T_m, \mathcal{D}_m)$ be the counterfactual version of $X = \min(T, \mathcal{D})$ had "the time m dietary intervention" been carried out. Then we make the following more realistic assumption.

Realistic Comparability (RC) Assumption : $A(m)$ is statistically independent of (Y_m, X_m) given $\Xi(m) = 1, \bar{L}(m), \bar{A}(m-1)$ and $\bar{U}(m) = \bar{0}(m)$, where $U(m) = 1$ if a subject has at m or had prior to m , an undiagnosed chronic disease that was sufficiently advanced to interfere with his normal weight trajectory. Otherwise $U(m) = 0$. We also define $U(m) = 1$ for subject's alive at m with $X < m$ under the assumption that there was probably a subclinical period prior to the time X of clinical diagnosis in which weight gain may have been altered. Note that $U(m) = 0$ implies $\bar{U}(m) = (U(0), \dots, U(m)) = \bar{0}(m)$ is also zero.

Remark: The RC assumption cannot be recast as (Y_m, X_m) independent of $A(m)$ given $(\bar{L}(m), \bar{A}(m-1), \bar{U}(m))$ even had we used the coding convention that vector $L(k)$ includes $\Xi(k)$ as a component, because, even under this coding, $pr(A(k) = 0 | \bar{A}(k-1) = \bar{0}(k-1), \bar{L}(k), \bar{U}(m), (Y_m, X_m))$ would be neither zero nor one and could, under the RC assumption, depend on (Y_m, X_m) whenever $\bar{U}(m) \neq 0$ and $\Xi(k) = 1$. For this reason it would perhaps be more precise to refer to the RC as a selective comparability assumption as it only implies comparability for a selected subset of the population.

We observe (Y, X, \bar{L}, \bar{A}) but $\bar{U}(m) = (U(0), \dots, U(m))$ is, of course, generally unobserved when $X > m$. Thus \bar{U} is an unmeasured confounder. The most crucial of several assumptions needed to allow consistent estimation of the parameter of interest $E[Y_0]$ in this setting is the following.

Clinical Detection (CD) Assumption: Any subject who has $U(m) = 1$ [ie a sufficiently advanced undiagnosed chronic disease at (or before) m] and

thereafter follows "the time m dietary intervention" will either have died or been diagnosed with clinical chronic disease by time $m + \zeta$, where ζ is assumed known. Formally

$$U(m) = 1 \Rightarrow X_m \leq m + \zeta \quad (14)$$

or equivalently

$$X_m > m + \zeta \Rightarrow \overline{U}(m) = \overline{0}(m) \quad (15)$$

Here \Rightarrow is translated as 'implies'. A typical choice for ζ might be 72 months. It is useful to choose ζ to be the minimal time for which (15) holds as this increases both the efficiency of g-estimators and the power of goodness of fit tests to detect misspecification of a structural nested model and decreases the likelihood that an SNM is nonidentified. However if the chosen ζ is less than the true minimum time for which (15) holds bias will result. As a consequence one should routinely include a table that shows how one's estimate of $E[Y_0]$ changes as ζ is varied.

The RC and CD assumptions require that one record in X the minimum time of clinical onset among the set of clinical conditions whose preclinical phase could affect BMI. The exact clinical conditions that belong in this subset is a substantive question, about which subject matter experts should be consulted.

Remark: We will later consider the effect of replacing the counterfactual X_m by the observed X in the CD assumption.

3.2.1 Estimation under a Rank Preserving SNM for $Y_m|X_m$ with X_m known

To consistently estimate $E[Y_0]$ under *RC* and *CD* we must replace our SNMM model with an additive SNMM model for $Y_m|X_m$ that also conditions on and allows effect modification by the counterfactual X_m . For pedagogic purposes in this subsection we return to locally rank preserving models. A locally rank preserving SNM for $Y_m|X_m$ states that

$$Y_{m+1} - Y_m = \gamma_m[A(m), \overline{A}(m-1), \overline{L}(m), X_m, \beta^*] \quad (16)$$

where β^* is an unknown parameter and $\gamma_m[A(m), \overline{A}(m-1), \overline{L}(m), X_m, \beta]$ is a known function that can now depend on X_m that takes the value zero if either $A(m) = 0$ or $\beta = 0$. [We emphasize that it is X_m and not \overline{X}_m that occurs in the last display.]. This model is equivalent to assuming

$$Y_m = Y_m(\beta^*) \quad (17)$$

for each subject with $Y_m(\beta)$ now redefined as

$$Y_m(\beta) = Y - \sum_{j=m}^K \gamma_m[A(j), \overline{A}(j-1), \overline{L}(j), X_j, \beta] \quad (18)$$

Now, of course the counterfactual variable X_m is itself unobserved. However for pedagogic purposes in this subsection we unrealistically assume that in

addition to the observed data (Y, X, \bar{L}, \bar{A}) , data on the counterfactuals X_m are available.

Remark: We do not actually require a locally rank preserving SNM for $Y_m|X_m$. A locally rank preserving SNM for $Y_m|c(X_m)$ for a certain known function $c(x)$ could be used instead. This remark is explored further in the Appendix.

Redefine $Q_m(\beta) = q_m[\bar{L}(m), \bar{A}(m-1), X_m, Y_m(\beta)]$ to possibly be a function of X_m . Consider again the g-estimator $\hat{\beta}$ that is equal to the β for which the 5 degree of freedom score test of $\theta = 0$ is precisely zero in the model

$$E[A(m) | \bar{L}(m), \bar{A}(m-1), X_m, \Xi(m) = 1, Y_m(\beta)] = \alpha^T W(m) + \theta^T Q_m(\beta).$$

$\hat{\beta}$ would be a CAN estimator of β^* under the CO assumption, but not under the RC assumption. Under RC, the independence needed to make $\theta = 0$ when $\beta = \beta^*$ only holds when we also condition on $\bar{U}(m)$.

However, consider the estimator $\tilde{\beta}$ obtained when, for each time m , we only fit the previous model to subjects for whom $X_m > m + \zeta$, excluding all subjects with $X_m \leq m + \zeta$. This exclusion can be expressed by saying that we now fit the model

$$E[A(m) | \bar{L}(m), \bar{A}(m-1), X_m, \Xi(m) = 1, Y_m(\beta), X_m > m + \zeta] = \alpha^T W(m) + \theta^T Q_m(\beta). \quad (19)$$

Then the estimator $\tilde{\beta}$ is the β for which the 5 degree of freedom score test of the hypothesis $\theta = 0$ is precisely zero in this latter model. When $X_m > m + \zeta$, $\bar{U}(m) = \bar{0}(m)$, by assumption CD. Hence we can rewrite the last display as

$$E[A(m) | \bar{L}(m), \bar{A}(m-1), X_m, \Xi(m) = 1, Y_m(\beta), X_m > m + \zeta, \bar{U}(m) = \bar{0}(m)] = \alpha^T W(m) + \theta^T Q_m(\beta) \quad (20)$$

showing that we have succeeded in conditioning on $\bar{U}(m) = \bar{0}(m)$, even though $\bar{U}(m)$ is unmeasured! It follows that, when the parameter β^* is identified, the estimator $\tilde{\beta}$ is a consistent and asymptotically normal (CAN) estimator of β^* under the RC and CD assumptions, since these assumptions imply the coefficient $\theta = 0$ if $\beta = \beta^*$. However as discussed further below, under the RC and CD assumptions, the positivity assumption no longer suffices to guarantee identification.

In summary, all that was required to produce a CAN estimator $\tilde{\beta}$ of the parameter β^* of our locally rank preserving SNM (17) for $Y_m|X_m$ under the RC and CD assumptions was to restrict the earlier g-estimation procedure at each time m to those subjects with $X_m > m + \zeta$.

Thus if $\gamma_m[A(m), \bar{A}(m-1), \bar{L}(m), X_m, \beta] = A(m) \beta^T R_m$ is linear in β with $R_m = r_m(\bar{L}(m), X_m)$ being a vector of known functions that now can depend on X_m , then, given the OLS estimator $\hat{\alpha}^T$ of α^T in the model $E[A(m) | \bar{L}(m), \bar{A}(m-1), \Xi(m) = 1] = \alpha^T W(m)$ and $Q_m(\beta) = Q_m^* Y_m(\beta)$ lin-

ear in $Y_m(\beta)$, the CAN estimator $\tilde{\beta}$ exists in closed form as

$$\tilde{\beta} = \left\{ \sum_{i=1, m=0}^{i=n, m=K} I[X_{im} > m + \zeta] \Xi_i(m) A_i(m) G_{im}(\hat{\alpha}) Q_{im}^* S_{im}^T \right\}^{-1} \quad (21)$$

$$\times \left\{ \sum_{i=1, m=0}^{i=n, m=K} I[X_{im} > m + \zeta] \Xi_i(m) Y_i G_{im}(\hat{\alpha}) Q_{im}^* \right\}$$

with $G_{im}(\hat{\alpha}) = [A_i(m) - \hat{\alpha}^T W_i(m)]$, $S_{im} = \sum_{i=1, j=m}^{i=n, j=K} R_{ij}$.

From the above, it follows that if, in addition to the observed data (Y, X, \bar{L}, \bar{A}) , data on the counterfactuals X_m are available for each m , the sample average $\sum_i^n Y_0(\tilde{\beta})/n$ is a CAN estimator of the parameter of interest $E[Y_0]$ under the RC and CD assumptions, provided β^* is identified and both our locally rank preserving SNM for $Y_m|X_m$ and our model

$$E[A(m) | \bar{L}(m), \bar{A}(m-1), \Xi(m) = 1] = \alpha^T W(m) \quad (22)$$

are correct. Of course data on X_m are unavailable. However in the next subsection we prove an analogue of this result holds without data on X_m under a locally rank preserving SNFTM for the X_m which allows us to replace X_m by an estimate $X_m(\tilde{\psi})$, where $\tilde{\psi}$ estimates the parameter ψ^* of our SNFTM.

Before preceding to the next subsection, several additional points need to be made.

Can we replace X_m by X : A natural question that arises is the following. Suppose we replaced X_m by the observed X in the CD assumption, in our definition of $Y_m(\beta)$, and wherever else X_m occurs in this subsection, with the exception of the RC assumption (as the RC assumption with X replacing X_m would clearly be false if BMI is a cause of T and/or C and thus of X). Do $\tilde{\beta}$ and $\sum_i^n Y_0(\tilde{\beta})/n$ remain CAN estimators of β^* and $E[Y_0]$? This question is natural in the sense that it is not obvious that the CD assumption and RP SNM based on X_m are more likely to be true than when based on X . So were the answer "yes" it would be simpler and more straightforward to use X in place of X_m . In particular, since X , unlike X_m , is observed, we would eliminate the need to replace X_m with the estimator $X_m(\tilde{\psi})$, thereby greatly simplifying the analysis.

Unfortunately $\tilde{\beta}$ and thus $\sum_i^n Y_0(\tilde{\beta})/n$ do not remain consistent when we use X in place of X_m . To see why consider the model

$$E[A(m) | \bar{L}(m), \Xi(m) = 1, \bar{A}(m-1), X, Y_m(\beta), X > m + \zeta] = \alpha^T W(m) + \theta^T Q_m Y_m(\beta) \quad (23)$$

which has replaced X_m in Eq (19) with X . Clearly $\tilde{\beta}$ will only be consistent for the parameter β^* of our locally RP SNM if $\theta = 0$ when $\beta = \beta^*$. That is $\tilde{\beta}$ will only be consistent if $Y_m = Y_m(\beta^*)$ is independent of $A(m)$ given

$(\bar{L}(m), \bar{A}(m-1), \Xi(m) = 1, X, X > m + \zeta)$. Now by the CD assumption with X replacing X_m , $X > m + \zeta$ implies $\bar{U}(m) = \bar{0}(m)$. Thus consistency of $\tilde{\beta}$ requires $Y_m(\beta^*)$ independent of $A(m)$ given

$(\bar{L}(m), \bar{A}(m-1), \Xi(m) = 1, X, X > m + \zeta, \bar{U}(m) = \bar{0}(m))$. However, we show in the next paragraph that this independence statement is not implied by the RC assumption and thus will generally be false, unless $A(k)$ has no causal effect on X for $k \geq m$ in which case $X = X_m$ for each subject and we are back to Eq. (20).

When $A(m)$ has a causal effect on X (whether directly or through $A(k)$, $k > m$) then X is a common effect of two causes $A(m)$ and X_m that are independent conditional on the event $(\bar{L}(m), \bar{A}(m-1), \Xi(m) = 1, \bar{U}(m) = \bar{0}(m), Y_m(\beta^*))$. Therefore, conditional on both the previous event and $(X, X > m + \zeta)$, $A(m)$ and X_m are dependent and thus so are $A(m)$ and $Y_m(\beta^*)$, since X_m and $Y_m(\beta^*)$ are highly correlated, as both are functions of T_m .

However, even when $A(m)$ has a causal effect on X , a slight modification of the above estimation procedure can be used to obtain CAN estimators of β^* in the special case in which $A(m)$ has a known minimal latent period χ for its effect on X of at least ζ months.

Definition of Minimal Latent Period (MLP) for effect on X : $A(m)$ has a minimal latent period for its effect on X of χ months if, for every subject and each time $k > m$, $X_k > m + \chi \Leftrightarrow X_m > m + \chi$ and $X_k = X_m$ if $X_m < m + \chi$. In particular by taking $k = K + 1$, the last two statements become $X > m + \chi \Leftrightarrow X_m > m + \chi$ and $X = X_m$ if $X_m < m + \chi$.

When a known minimal latent period χ exceeds ζ , we can obtain CAN estimators of β^* by simply replacing $X > m + \zeta$ by $m + \chi > X > m + \zeta$ in model (23) since then the event $(X, m + \chi > X > m + \zeta)$ is the event $(X_m, m + \chi > X_m > m + \zeta)$ and we are back in the setting of Eq. (19), except for the additional restriction, $m + \chi > X_m$, which does not introduce bias. Thus the existence of a minimal latent period of length χ greater than ζ allows us to estimate β^* and $E[Y_0]$ without the need to specify a SNFTM for the X_m .

We now prove that under the RC and CD assumptions, a MLP of length χ greater than ζ implies that

$$X \perp\!\!\!\perp A(m) \mid \bar{L}(m), \Xi(m) = 1, \bar{A}(m-1), m + \chi > X > m + \zeta.$$

It follows that taking the RC and CD assumptions as given, we can test the hypothesis that a MLP of length χ greater than ζ exists by testing whether the last display is true. In fact a test of the hypothesis that the parameter ψ^* of a SNFTM serves as a test of the previous display. To prove our previous claim note that, by the MLP assumption, the event $m + \chi > X > m + \zeta$ is the event $m + \chi > X_m > m + \zeta$ which, by the CD assumption, is the event $m + \chi > X_m > m + \zeta, \bar{U}(m) = \bar{0}(m)$. Thus the last display is under the MLP and CD assumption equivalent to the statement " X_m is independent of $A(m)$ given $(\bar{L}(m), \Xi(m) = 1, \bar{A}(m-1), m + \chi > X_m > m + \zeta, \bar{U}(m) = \bar{0}(m))$ " which is true by the RC assumption.

Most experts believe it to be substantively implausible that an increase in

BMI has a minimum latent period of more than 72 months, our default choice for ζ . In contrast, in occupational cohort studies of the effect of a chemical carcinogen on time to clinical cancer, minimum latent periods of up to 10 years are commonly assumed.

3.2.2 Estimation of $E[Y_0]$ under a Rank Preserving SNFTM:

As mentioned above, an analogue of the above results hold when data on X_m are unavailable under a locally rank preserving SNFTM for X_m . The simplest locally rank preserving SNFTM specifies that

$$X_m = m + \int_m^X \exp(\psi^* A(t)) dt \text{ if } X > m \quad (24)$$

$$X_m = X \text{ if } X \leq m, \quad (25)$$

where ψ^* is an unknown parameter and $A(t)$ is as defined previously when t is a whole number of months and $A(t) = A(\lfloor t \rfloor)$ when t is not a whole number where $\lfloor t \rfloor$ is the largest integer less than or equal to t . Thus, by the definition of an integral as the area under a curve,

$$\int_m^X \exp(\psi^* A(t)) dt = \sum_{j=m}^{j=\lfloor X-1 \rfloor} \exp(\psi^* A(j)) + \{X - \lfloor X \rfloor\} \exp(\psi^* A(\lfloor X \rfloor)).$$

A locally rank preserving SNFTM directly maps an individual's observed failure time X to the failure time X_m the individual would have under the "time m dietary intervention". Thus it is a model for individual causal effects. If $\psi^* = 0$, $\exp(\psi^* A(t)) = 1$ and thus $X_m = m + \int_m^X dt = m + X - m = X$ for any m . Hence $\psi^* = 0$ encodes the sharp null hypothesis that $X_0 = X$ for all subjects, i.e., the "time 0 dietary intervention" has no effect on any subject's $X = \min(T, \mathcal{D})$. It is useful to note that when $\psi^* \neq 0$, the SNFTM (24)-(25) implies that there is no minimal latent period for the effect of treatment on X .

A general class (although not the most general class) of locally RP SNFTMs that includes the above one parameter model assumes

$$X_m = m + \int_m^X \exp\{\omega(\bar{A}(t), \bar{L}(t), \psi^*)\} dt \text{ if } X > m \quad (26)$$

$$X_m = X \text{ if } X \leq m \quad (27)$$

where $\omega(\bar{A}(t), \bar{L}(t), \psi) \equiv \omega(A(t), \bar{A}(t^-), \bar{L}(t), \psi)$ is a known function satisfying $\omega(A(t), \bar{A}(t^-), \bar{L}(t), \psi) = 0$ if $A(t) = 0$ or $\psi = 0$ and $\bar{A}(t^-)$ is the A-history until just prior to time t . For example, we might have $\omega(A(t), \bar{A}(t^-), \bar{L}(t), \psi) = A(t) \{\psi_0 + \psi_1^T L(t)\}$ where $L(t) = L(\lfloor t \rfloor)$ and $L(\lfloor t \rfloor)$ is as defined earlier.

We next turn to estimation of ψ^* . For the moment, suppose the CO assumption modified to have (Y_m, X_m) in place of Y_m held and that (Y, X, \bar{L}, \bar{A}) was

observed. Then we could consistently estimate ψ^* by g-estimation. Specifically we define

$$X_m(\psi) = m + \int_m^X \exp\{\omega(\bar{A}(t), \bar{L}(t), \psi)\} dt \text{ if } X > m \quad (28)$$

$$X_m(\psi) = X \text{ if } X \leq m \quad (29)$$

so under our model, $X_m = X_m(\psi^*)$. Note that, for each ψ , $X_m(\psi)$ can be computed from the observed data. Suppose, for concreteness, ψ^* is 5 dimensional so we search over a 5 dimensional grid. We let $Q_m^{**}(\psi) = q_m^{**}[\bar{L}(m), \bar{A}(m-1), X_m(\psi)]$ be a 5 dimensional vector of functions of $(\bar{L}(m), \bar{A}(m-1), X_m(\psi))$ such as $Q_m^{**}(\psi) = X_m(\psi)[1, m, L^T(m)]$. We use an extended linear model

$$E[A(m) | \bar{L}(m), \bar{A}(m-1), \Xi(m) = 1, X_m(\psi)] = \alpha^T W(m) + \theta^T Q_m^{**}(\psi)$$

Our g-estimate $\hat{\psi}$ is the ψ for which the 5 degree of freedom score test that all 5 components of θ equal zero is precisely zero. Since, by the modified CO assumption, $\theta = 0$ if $\psi = \psi^*$, the g-estimate $\hat{\psi}$ is CAN for ψ^* . The particular choice of the functions $Q_m^{**}(\psi)$ does not affect the consistency of the point estimate, but it determines the width of its confidence interval. Because $X_m(\psi)$ is a nonlinear function of ψ , there is not a closed form expression for $\hat{\psi}$. However the equation solved by $\hat{\psi}$ is a smooth function of ψ , so standard methods for solving nonlinear equations such as the Newton-Raphson algorithm can be used to compute $\hat{\psi}$.

Next suppose the observed data is still (Y, X, \bar{L}, \bar{A}) , but the modified CO assumption does not hold. Rather, the CD and RC assumptions hold. Define the estimator $\tilde{\psi}$ as the ψ for which the 5 degree of freedom score test of the hypothesis $\theta = 0$ is precisely zero in the model

$$E[A(m) | \bar{L}(m), \bar{A}(m-1), \Xi(m) = 1, X_m(\psi), X_m(\psi) > m + \zeta] = \alpha^T W(m) + \theta^T Q_m^{**}(\psi).$$

Note the set of subjects who do not contribute to the score test of $\theta = 0$ (i.e subjects with $X_m(\psi) \leq m + \zeta$) depends on ψ . When $X_m = X_m(\psi^*) > m + \zeta$, then $\bar{U}(m) = \bar{0}(m)$, by assumption CD. Hence, at $\psi = \psi^*$, our procedure conditions on $\bar{U}(m) = \bar{0}(m)$. It follows that, provided ψ^* is identified, the estimator $\tilde{\psi}$ is a CAN estimator of ψ^* under the RC assumption, as that assumption implies the coefficient $\theta = 0$ if $\psi = \psi^*$. However, under the CD and RC assumptions, the positivity assumption does not guarantee identification.

Now let $\tilde{\beta}(\tilde{\psi})$ be defined like $\tilde{\beta}$ except that everywhere $X_m(\tilde{\psi})$ replaces X_m , so that $\tilde{\beta}(\tilde{\psi})$ is a function of the data (Y, X, \bar{L}, \bar{A}) only. Next define

$$Y_m(\beta, \psi) = Y - \sum_{j=m}^K \gamma_m[A(j), \bar{A}(j-1), \bar{L}(j), X_j(\psi), \beta] \quad (30)$$

Note, by both models (17) and (26-27) being locally rank preserving, $Y_m(\beta^*, \psi^*) = Y_m$. Thus, when ψ^* and β^* are identified, the sample average $\sum_i^n Y_0 \left[\left(\tilde{\beta}(\tilde{\psi}), \tilde{\psi} \right) \right] / n$ is a CAN estimator of the parameter of interest $E[Y_0]$ under the RC and CD assumptions, provided both our locally rank preserving SNM (17) for $Y_m|X_m$, our locally rank preserving SNFTM (26-27) for X_m , and our model (22) are all correctly specified.

3.2.3 Estimation of $E[Y_0]$ under a SNMM and a SNFTM without Rank Preservation:

As discussed earlier, the assumption of local rank preservation is biologically implausible. Thus we will no longer assume that our locally rank preserving models (17) and (26-27) are true. As a consequence we can no longer assume that there exists some (β^*, ψ^*) such that the unobserved counterfactuals (X_m, Y_m) equal the observed $(X_m(\psi), Y_m(\beta, \psi))$ when $(\beta, \psi) = (\beta^*, \psi^*)$. However, suppose with $(X_m(\psi), Y_m(\beta, \psi))$ still defined by (28-29), and (30), we assume that, for each m , there exists some (β^*, ψ^*) such that

Assumption (i): when $\psi = \psi^*$, X_m and $X_m(\psi)$ have the same conditional distribution given $(A(m), \bar{L}(m), \bar{A}(m-1))$ and

Assumption (ii)

$$E[Y_m|A(m), \bar{L}(m), \bar{A}(m-1), X_m = x] = E[Y_m(\beta^*, \psi^*)|A(m), \bar{L}(m), \bar{A}(m-1), X_m(\psi^*) = x] \quad (31)$$

In contrast with the assumption of local RP, there is no apriori biological reason to exclude the possibility that (i) and (ii) both hold.

When assumptions (i) and (ii) hold for each m , we say the SNMM

$$\gamma_m[A(m), \bar{A}(m-1), \bar{L}(m), x, \beta] \quad (32)$$

for $Y_m|X_m$ and the SNFTM (28) – (29) for X_m jointly hold with true parameter (β^*, ψ^*) . If the RC and CD assumptions, the model (22) and (i) and (ii) all hold, then $\tilde{\psi}, \tilde{\beta}(\tilde{\psi})$, and $\sum_i^n Y_0 \left[\left(\tilde{\beta}(\tilde{\psi}), \tilde{\psi} \right) \right] / n$ as defined previously are CAN for ψ^*, β^* , and the parameter of interest $E[Y_0]$ respectively, provided (β^*, ψ^*) are identified and we choose $Q_m(\beta)$ linear in $Y_m(\beta)$. [In contrast, $Q_m^*(\psi)$ need not be chosen linear in $X_m(\psi)$]. In summary, $\tilde{\psi}, \tilde{\beta}(\tilde{\psi})$, and $\sum_i^n Y_0 \left[\left(\tilde{\beta}(\tilde{\psi}), \tilde{\psi} \right) \right] / n$ have the same statistical properties under our joint SNMM model for $Y_m|X_m$ and SNFTM for X_m when local rank preservation does not hold as when it does.

3.2.4 Are Remarkable Results due to Some Sleight of Hand

The result summarized in the last sentence is striking for a number of reasons. Our comparability assumption, i.e. the RC assumption, only assumes no unmeasured confounding conditional on $\bar{U}(m)$. Yet neither the SNMM for $Y_m|X_m$ nor SNFTM for X_m is a model for causal effects conditional on the unmeasured

$\overline{U}(m)$. Thus, it is remarkable that these models can be used to estimate causal contrasts such as $E[Y_0] - E[Y]$ under the RC and CD assumptions. Furthermore, even though $X_m > m + \zeta$ implies $\overline{U}(m) = \overline{0}(m)$ by the CD assumption, nonetheless, in the absence of local rank preservation, $X_m(\psi^*) > m + \zeta$ does not imply $\overline{U}(m) = \overline{0}(m)$. Hence when local rank preservation does not hold, even though we condition on $X_m(\psi) > m + \zeta$ in computing our g-estimates $\tilde{\psi}, \tilde{\beta}(\tilde{\psi})$, we do not thereby restrict the analysis to a subset of subjects all of whom have the same value of $\overline{U}(m)$; thus one might guess confounding by the unmeasured $\overline{U}(m)$ has not been controlled and our estimates of $\tilde{\psi}, \tilde{\beta}(\tilde{\psi})$, and $\sum_{i=1}^n Y_{0,i} \left[\left(\tilde{\beta}(\tilde{\psi}), \tilde{\psi} \right) \right] / n$ must be inconsistent. Remarkably, such is not the case.

How did we pull off the seemingly remarkable ‘magic’ described in the preceding paragraph? We shall investigate whether we used some subtle “sleight of hand”. We use a simple paradigmatic instance of our model that only involves a single time-independent exposure to guide our investigation. Specifically, we next provide an explicit proof that contains no “sleight of hand” of our results in the case of a time-independent exposure. The general case is treated in the appendix. The reader who is interested more in the methodology and less interested in foundational issues may feel free to skip ahead to section 4.

Paradigmatic Instance of a Time-Independent Exposure: We suppose that $K + 1 = 1$ so time 0 is the only time of exposure. Further we assume there are no covariates. In this setting the RC assumption becomes (Y_0, X_0) independent of $A(0)$ given the unmeasured confounder $U(0) = 0$. The CD assumption becomes $X_0 > \zeta$ implies $U(0) = 0$. Our SNFTM for X_0 becomes

Assumption (i): $X_0(\psi) = X \exp(\psi A(0))$ and X_0 have the same conditional distribution given $A(0)$ at $\psi = \psi^*$,

while our SNMM for $Y_m | X_m$ becomes

Assumption (ii): $E[Y_0 | A(0), X_0 = x] = E[Y_0(\beta^*, \psi^*) | A(0), X_0(\psi^*) = x]$ where $Y_0(\beta, \psi) = Y - \gamma_0[A(0), X_0(\psi), \beta]$.

Neither model makes any reference to $U(0)$ and thus neither is a model for causal effects conditional on $U(0)$. Furthermore, although $X_0 > \zeta$ implies $U(0) = 0$ by the CD assumption, nonetheless $X_0(\psi^*) > \zeta$ does not imply $U(0) = 0$. Now to prove our results.

Proofs Of Our Results: Proof that $\tilde{\psi}$ is CAN for ψ^* : By assumption (i), $pr[X_0(\psi^*) > t | A(0), X_0(\psi^*) > \zeta] = pr[X_0 > t | A(0), X_0 > \zeta]$. But, by the CD and then the RC assumptions, $pr[X_0 > t | A(0), X_0 > \zeta] = pr[X_0 > t | A(0), X_0 > \zeta, U(0) = 0] = pr[X_0 > t | X_0 > \zeta, U(0) = 0]$. Hence $pr[X_0(\psi^*) > t | A(0), X_0(\psi^*) > \zeta]$ is not a function of $A(0)$. We conclude that $A(0)$ and $X_0(\psi^*)$ are independent given $X_0(\psi^*) > \zeta$. Thus $E[A(0) | X_0(\psi^*), X_0(\psi^*) > \zeta] = \alpha + \theta X_0(\psi^*)$ has coefficient $\theta = 0$ so, when ψ^* is identified, the $\tilde{\psi}$ for which the the score test of $\theta = 0$ takes the value 0 is CAN for ψ^* .

Proof that $\tilde{\beta}(\tilde{\psi})$ is CAN for β^* : By Assumption (ii),

$$E[Y_0(\beta^*, \psi^*) | A(0), X_0(\psi^*) = x, X_0(\psi^*) > \zeta] = E[Y_0 | A(0), X_0 = x, X_0 > \zeta].$$

But, by the CD and then the RC assumptions,

$$E[Y_0 | A(0), X_0 = x, X_0 > \zeta] = E[Y_0 | A(0), X_0 = x, X_0 > \zeta, U(0) = 0]$$

$$= E[Y_0 | X_0 = x, X_0 > \zeta, U(0) = 0] \text{ is not a function of } A(0).$$

$$\text{Thus, } 0 = E[Y_0(\beta^*, \psi^*) \{A(0) - E[A(0) | X_0(\psi^*), X_0(\psi^*) > \zeta]\}].$$

Hence $0 = E[Y_0(\beta^*, \psi^*) \{A(0) - E[A(0) | X_0(\psi^*) > \zeta]\}]$. As a consequence, the $\tilde{\beta}(\tilde{\psi})$ for which the the score test of $\theta = 0$ takes the value 0 in the model

$$E[A(0) | X_0(\tilde{\psi}) > \zeta, Y_0(\beta, \tilde{\psi})] = \alpha + \theta Y_0(\beta, \tilde{\psi}) \text{ is CAN for } \beta^*, \text{ when } \beta^* \text{ and } \psi^* \text{ are identified.}$$

Proof that $\sum_{i=1}^n Y_{0,i} [\tilde{\beta}(\tilde{\psi}), \tilde{\psi}] / n$ is CAN for $E[Y_0] : E[Y_0] =$

$$\int \int E[Y_0 | A(0), X_0 = x] dF_{X_0}(x | A_0) dF(A_0)$$

$$= \int \int E[Y_0(\beta^*, \psi^*) | A(0), X_0(\psi^*) = x] dF_{X_0(\psi^*)}(x | A_0) dF(A_0) = E[Y_0(\beta^*, \psi^*)]$$
 by assumptions (i) and (ii). Hence, $\sum_{i=1}^n Y_{0,i} [\tilde{\beta}(\tilde{\psi}), \tilde{\psi}] / n$ is CAN for $E[Y_0(\beta^*, \psi^*)] = E[Y_0]$, when β^* and ψ^* are identified.

This completes the promised proof of our results in the time-independent case. The proof in the appendix of the general time-dependent case is not much more difficult when one proceeds by induction. We conclude no sleight of hand occurred in the proof.

Do Correctly Specified SNMMs for $Y_m | X_m$ and SNFTMs for X_m Always Exist? Perhaps the sleight of hand occurred right at the start, when we supposed that there exist (β^*, ψ^*) such that assumptions (i) and (ii) hold. We now prove that no such sleight of hand is afoot. Specifically we prove that there always exist correctly specified SNMMs for $Y_m | X_m$ and SNFTMs for X_m . [This result does not, of course, imply that the particular SNMM and SNFTM we actually choose to analyze are correct.] We actually prove this result for an alternative, more intuitive, definition of a SNMM for $Y_m | X_m$ and a SNFTM for X_m and then prove these alternative definitions are logically equivalent to assumptions (i) and (ii). This is done in this subsection for the special case of a time-independent exposure and in the Appendix for a general time-varying exposure.

Consider again, for simplicity, our paradigmatic instance. Write $A(0)$ as A . Suppose that Y, Y_0, X, X_0 are all non-negative continuous random variables with support on $(0, \infty)$, satisfying the consistency assumption $X = X_0$ and $Y = Y_0$ if $A = 0$. Let $S(x|A) = \text{pr}(X > x | A)$. Let $S_0(x|A) = \text{pr}(X_0 > x | A)$. Let $S_0^{-1}(x|A)$ be the inverse of $S_0(x|A)$ wrt the x argument. Define the function $x_0^\dagger(x, A) = S_0^{-1}[\{S(x|A)\} | A]$. Substituting 0 for A , we find $x_0^\dagger(x, 0) = x$, so

$$x_0^\dagger(X, 0) = X \text{ wp1} \tag{33}$$

Define $X_0^\dagger = x_0^\dagger(X, A)$. Then $X_0^\dagger = x_0^\dagger(X, 0) = X$, when $A = 0$. It is well known that $X_0^\dagger = x_0^\dagger(X, A)$ and X_0 have the same conditional distribution given A .

Define $S(t|A, X_0^\dagger = x) = \text{pr}(Y > t|A, X_0^\dagger = x)$ and $S_0(t|A, X_0 = x) = \text{pr}(Y_0 > t|A, X_0 = x)$. Let $S_0^{-1}(t|A, X_0 = x)$ be the inverse of $S_0(t|A, X_0 = x)$ wrt the t argument. Let $y_0^\dagger(t, x, A) = S_0^{-1}\left(\left\{S(t|A, X_0^\dagger = x)\right\}|A, X_0^\dagger = x\right)$ and $Y_0^\dagger = y_0^\dagger(Y, X, A)$. Then $Y_0^\dagger|A, X_0^\dagger = x$ and $Y_0|A, X_0 = x$ have the same conditional distribution. It follows that $(Y_0^\dagger, X_0^\dagger)|A$ and $(Y_0, X_0)|A$ have the same joint conditional distribution. Thus,

$$E\left[Y_0^\dagger|X_0^\dagger = x, A\right] = E[Y_0|X_0 = x, A] \quad (34)$$

Define

$$\gamma^\dagger(A, x) = E\left[Y|X_0^\dagger = x, A\right] - E\left[Y_0^\dagger|X_0^\dagger = x, A\right] \equiv E\left[Y - Y_0^\dagger|X_0^\dagger = x, A\right] \quad (35)$$

The last two displays imply that

$$E\left[Y - \gamma^\dagger(A, X_0^\dagger)|X_0^\dagger = x, A\right] = E[Y_0|X_0 = x, A] \quad (36)$$

and

$$\gamma^\dagger(0, X) = 0 \text{ wp1} \quad (37)$$

since, by $Y_0^\dagger = y_0^\dagger(Y, X, A)$, $Y_0^\dagger = Y$ when $A = 0$.

Here are the alternative definitions of a SNFTM for X_0 and a SNMM for $Y_0|X_0$.

Definition a: Let $x_0(t, a, \psi)$ be known function montone increasing in t for each (a, ψ) satisfying $x_0(t, a, \psi) = 1$ if $a = 0$ or $\psi = 0$. We say $x_0(t, a, \psi)$ is a correctly specified SNFTM for X_0 if there exists ψ^* such that $X_0(\psi^*) \equiv x_0(X, A, \psi^*)$ equals X_0^\dagger with probability one.

Definition b: We say a known function $\gamma(a, x, \beta)$ satisfying $\gamma(a, x, \beta) = 0$ if $a = 0$ or $\beta = 0$ is a correctly specified SNMM for $Y_0|X_0$ if, for some β^* , $\gamma(A, X, \beta^*) = \gamma^\dagger(A, X)$ with probability 1.

Define $Y_0(\beta^*, \psi^*) = Y - \gamma(A, X_0(\psi^*), \beta^*)$.

It is obvious from definitions a and b that there always exist correctly specified SNMMs for $Y_0|X_0$ under Definition b and correctly specified SNFTMs for X_0 under definition a since $\gamma^\dagger(A, X)$ and $x_0^\dagger(X, A)$ are well defined functions of (F, F_0) satisfying $\gamma^\dagger(0, X) = 0$ and $x_0^\dagger(X, 0) = X$ with probability one, where F and F_0 , respectively, denote the joint distribution of (Y, X, A) and of (Y_0, X_0, A) . Note $\gamma^\dagger(A, X)$ and $x_0^\dagger(X, A)$ do not depend on the conditional joint distribution of $\{(Y, X), (Y_0, X_0)\}$ given A . This is as desired as this joint is not non-parametrically identified from data (Y, X, A) even when A is randomly assigned.

Thus it only remain to show the logical equivalence of the original and alternative definitions of a SNFTM for X_0 and a SNMM for $Y_0|X_0$.

The following Lemma shows that the alternative definitions of a SNFTM for X_0 and a SNMM for $Y_0|X_0$ imply the previous definitions.

Lemma: Suppose $x_0(t, a, \psi)$ is a correctly specified SNFTM for X_0 as defined in definition **a**. Then $X_0(\psi^*)|A$ has the same distribution as $X_0|A$. Further assume that $\gamma(a, x, \beta)$ is a correctly specified SNMM for $Y_0|X_0$ as defined in definition **b**. Then $E[Y_0(\beta^*, \psi^*)|A, X_0(\psi^*) = x] = E[Y_0|X_0 = x, A]$.

Proof: The first result follows immediately from X_0^\dagger and X_0 having the same conditional distribution given A . The second result follows from

$$E[Y - \gamma^\dagger(A, x)|X_0^\dagger = x, A] = E[Y_0|X_0 = x, A].$$

Finally, the following Lemma shows that the original definitions imply the alternative definitions.

Lemma: Suppose $x_0(t, a, \psi)$ is montone increasing in t for each (a, ψ) satisfying $x_0(t, a, \psi) = 1$ if $a = 0$ or $\psi = 0$. Further suppose that $X_0(\psi^*)|A$ has the same distribution as $X_0|A$ wp1 where $X_0(\psi) = x_0(X, A, \psi)$. Then $X_0(\psi)$ is a correctly specified SNFTM for X_0 under definition **a**. In addition, suppose that $\gamma(a, x, \beta)$ is a function satisfying $\gamma(a, x, \beta) = 0$ if $a = 0$ or $\beta = 0$. Suppose $E[Y - \gamma(A, x, \beta^*)|X^\dagger = x, A = a] = E[Y_0|X_0 = x, A = a]$ for all (x, a) in a set of probability 1 under the law of (X_0, A) . Then, $\gamma(a, x, \beta)$ is a correctly specified SNMM for $Y_0|X_0$ under definition **b**.

Proof: The proof of the first part follows from the well known result that $X_0^\dagger = x_0^\dagger(X, A)$ is the only function $h(X, A)$ of (X, A) satisfying $h(X, A)|A$ has the same distribution as $X_0|A$ wp1. The second part is proved by showing that $\gamma^\dagger(a, x)$ is the unique function $h(a, x)$ satisfying $E[Y - h(A, x)|X^\dagger = x, A = a] = E[Y_0|X_0 = x, A = a]$ for all (x, a) in a set of probability 1 as in Refs (8,10).

Are $\gamma^\dagger(a, x), x_0^\dagger(x, a)$, and $E[Y_0]$ nonparametrically identified from data (Y, X, A) under our assumptions? In this subsection, we finally uncover some slight of hand that provided us with such seemingly magical results. Although we restrict our discussion to the special case of a time-independent exposure, similiar results apply in the general case. Specifically, we will show that $\gamma^\dagger(a, x), x_0^\dagger(x, a)$, and $E[Y_0]$ are not identified by the distribution of (Y, X, A) under the RC and CD assumptions. Previously, we saw that $\gamma^\dagger(a, x), x_0^\dagger(x, a)$, and $E[Y_0]$ are identified and equal $\gamma(a, x, \beta^*), x_0(x, a, \psi^*)$, and $E[Y - \gamma(A, X, \beta^*)]$, respectively when we assume a correctly specified SNFTM $x_0(x, a, \psi)$ for X_0 and a SNMM $\gamma(a, x, \beta)$ for $Y_0|X_0$ whose true parameters ψ^* and β^* are identified (by g-estimation). It follows that identification of $\gamma^\dagger(a, x), x_0^\dagger(x, a)$, and $E[Y_0]$ must result from the functional form restrictions encoded in our models $x_0(x, a, \psi)$ and $\gamma(a, x, \beta)$. It follows that if we make the restrictions imposed by our models less rigid by adding additional parameters, we can lose identification of $\gamma^\dagger(a, x), x_0^\dagger(x, a)$, and $E[Y_0]$. This loss of identification occurs when, in an infinite sample size, more than one combination of parameters, say the true parameters (ψ^*, β^*) and the false parameters (ψ^{**}, β^{**}) , both make the score

tests in our g-estimation procedures exactly zero for all choices of $Q_m(\beta)$ linear in $Y_m(\beta)$ and all choices of $Q_m^{**}(\beta)$. This loss of identification can be expressed by saying that the data (even were the sample size infinite) can not be used to determine whether the true causal quantities are $\gamma(a, x, \beta^*)$, $x_0(x, a, \psi^*)$, and $E[Y - \gamma(A, X, \beta^*)]$ versus $\gamma(a, x, \beta^{**})$, $x_0(x, a, \psi^{**})$, and $E[Y - \gamma(A, X, \beta^{**})]$.

In contrast, $\gamma^\dagger(a, x)$, $x_0^\dagger(x, a)$, and $E[Y_0]$ are identified under the comparability assumption that (Y_0, X_0) is independent of A_0 , without any reliance on the functional form restrictions encoded in our models. However, in contrast with assumption RC, this comparability assumptions contradicts our substantive knowledge, as it implies no unmeasured confounding by undiagnosed chronic disease.

The problem of lack of identification under the RC and CD assumptions has little to do with the question of local rank preservation. Suppose we have assumed a correctly specified SNFTM $x_0(x, a, \psi)$ for X_0 and we do not assume RP. Suppose in truth RP holds. Nonetheless, a second investigator who assumes the RP version of the SNFTM model gains nothing thereby in regard to the estimation of $x_0^\dagger(x, a)$: the causal quantity $x_0^\dagger(x, a)$ is identified under the non-rank preserving SNFTM if and only it is identified under the RP SNFTM. However, a small amount could be gained by assuming rank preservation for a SNMM; rarely by assuming RP a non-identifiable SNMM can become identifiable as one can then use non-linear functions $Q_m(\beta)$ of $Y_m(\beta)$ in g-estimation. But this advantage is not actually due to rank preservation. Rather it is due to the fact that an RP SNMM is actually a special case of a structural nested distribution model (SNDM) as defined in Refs (5) and (7). Our model SNMM model $\gamma(a, x, \beta)$ for $Y_0|X_0$ is a SNDM if $Y - \gamma(A, X, \beta^*)$ is independent (rather than just mean independent) of A given X . It is this independence (rather than rank preservation) that licences the use of non-linear functions $Q_m(\beta)$ of $Y_m(\beta)$ in g-estimation.

Non Identifiability of $\gamma^\dagger(a, x)$, $x_0^\dagger(x, a)$, and $E[Y_0]$ Suppose we do not impose a SNFTM for X_0 or a SNMM for $Y_0|X_0$. Then, it is clear that all we can conclude under assumptions RC and CD is that $X_0^\dagger = x_0^\dagger(X, A)$ and $A \equiv A(0)$ are independent given $X_0^\dagger > \zeta$ and $E[Y - \gamma_0^\dagger(A, x) | A(0), X_0^\dagger = x, X_0^\dagger > \zeta] = E[Y - \gamma_0^\dagger(A, x) | X_0^\dagger = x, X_0^\dagger > \zeta]$. As a consequence, our parameter of interest $E[Y_0]$ is not identified. Specifically, under RC and CD, with $p = pr(A = 0)$

$$E[Y_0] = \tag{38}$$

$$E[Y|X > \zeta, A = 0] \{pr[X > \zeta | A = 0]p + \{1 - pr[X^\dagger < \zeta | A \neq 0]\}\} (1 - p) \tag{39}$$

$$+ E[Y|X \leq \zeta, A = 0]pr[X \leq \zeta | A = 0]p \tag{40}$$

$$+ E\left[\left\{Y - \gamma^\dagger(A, X^\dagger)\right\} | X^\dagger \leq \zeta, A \neq 0\right]pr[X^\dagger < \zeta | A \neq 0](1 - p). \tag{41}$$

However the quantities

$$pr [X^\dagger < \zeta | A \neq 0] = pr [X_0 < \zeta | A \neq 0], \quad (42)$$

$$E \left[\left\{ Y - \gamma^\dagger (A, X^\dagger) \right\} | X^\dagger \leq \zeta, A \neq 0 \right] = E [Y_0 | X_0 \leq \zeta, A \neq 0] \quad (43)$$

are not identified under the RC and CD assumptions. It suffices to show this when RP holds. So, for the moment assume RP. Because both quantities (42) and (43) refer to the distribution of the counterfactuals responses (Y_0, X_0) under no exposure (no weight gain) among those who actually were exposed ($A \neq 0$), we need an assumption to identify them under RP. But under RC, we only have comparability conditional on a value of $U(0)$, which is unknown when $X_0 < \zeta$, so identification fails.

When we additionally assume a SNFTM for X_0 and a *SNMM* for $Y_0 | X_0$, we may or may not obtain identification of $E[Y_0]$ depending on whether the additional functional form restrictions encoded in the models suffice to identify the quantities (42) and (43) by allowing us to extrapolate from $X_0 > \zeta$ where we have comparability (since, by CD, $U(0) = 0$) to $X \leq \zeta$ where we do not. To clarify this last statement, consider the following RP SNM for $Y_0 | X_0 : Y_0 = Y - \gamma(A, X_0, \beta^*)$ with

$$\gamma(A, X_0, \beta) = \beta_0 AI(X_0 \leq \zeta) + \beta_1 AI(X_0 > \zeta). \quad (44)$$

Under assumptions RC and CD, even if we unrealistically suppose that data on X_0 was available for all subjects, we could not identify $\beta^* = (\beta_0^*, \beta_1^*)^T$, because β_0^* would not be identified, although β_1^* would be identified. This follows from the fact that, under RC and CD, no subject with $X_0 \leq \zeta$ may contribute to g-estimation of β^* . As a consequence we cannot identify $E[Y_0]$ because $Y_0 = Y - \beta_0^*$ is not estimable on the subset of exposed subjects ($A = 1$) with $X_0 \leq \zeta$.

In contrast, were data on X_0 available, β^* and $E[Y_0]$ are identified in the RP SNM $\gamma(A, X_0, \beta) = \beta_0 A + \beta_1 A X_0$ because both β_0^* and β_1^* can be estimated by g-estimation restricted to subjects with $X_0 > \zeta$. Thus $Y_0 = Y - \beta_0^* A - \beta_1^* A X_0$ can be estimated for all subjects, including those with $A = 1$ and $X_0 \leq \zeta$, because, by having the same parameters apply to subjects with $X_0 \leq \zeta$ as to subjects with $X_0 > \zeta$, the model allows extrapolation from subjects with $X_0 > \zeta$ to subjects with $X_0 \leq \zeta$. One must weigh the benefit of extrapolation that comes with assuming model $\gamma(A, X_0, \beta) = \beta_0 A + \beta_1 A X_0$ against the risk that the model is misspecified for subjects with $X_0 \leq \zeta$, as would be the case were the true model: $\gamma(A, X_0, \beta^*) = \beta_0^* AI(X_0 > \zeta) + \beta_1^* A X_0 I(X_0 > \zeta) + \beta_2^* AI(X_0 \leq \zeta) + \beta_3^* A X_0 I(X_0 \leq \zeta)$ with β_2^* very different from β_0^* and with β_3^* very different from β_1^* . Then the extrapolated value $Y - \beta_0^* A - \beta_1^* A X_0$ for Y_0 based on the misspecified model would be a badly biased estimate of the true Y_0 for subjects with $A = 1$ and $X_0 \leq \zeta$. Yet, because the model $\gamma(A, X_0, \beta) = \beta_0 A + \beta_1 A X_0$ is correct for subjects with $X_0 > \zeta$, there exists no valid test of model fit that could detect the biased extrapolation when we only assume RC and CD.

Suppose now, as is true in practice, data on X_0 are unavailable for subjects with $A = 1$. Then, under assumptions, RC and CD, without the help of a correct

RP SNFTM for X_0 whose functional form provides for extrapolation, we can no longer identify any aspect of the distribution of Y_0 for any identifiable subset of subjects with $A \neq 0$. This is because, although we know that the identified quantity $E[Y|X > \zeta, A = 0]$ equals $E[Y_0|X_0 > \zeta, A \neq 0]$, we cannot identify which subjects with $A \neq 0$ have $X_0 > \zeta$.

In summary, in the realistic setting of longitudinal time -dependent exposures, the possibility of sensitivity of one's estimate of $E[Y_0]$ to model extrapolation should be examined by reestimating $E[Y_0]$ under a variety of models that differ in both the dimension of the parameter vectors and in functional form.

A final point is that no individual who has developed a chronic disease by time m is included in our g-estimation procedure at m because $X_m(\psi) = X < m + \varsigma$ for such subjects. Thus our estimate of the effect of exposure at time m on a subject with a chronic disease at m is identified wholly by extrapolation from the effect on subjects without chronic disease at m . One approach to lessening the degree of extrapolation is to require a subject to be rather ill before they meet the definition of having a diagnosed chronic disease. For example, mild to moderate diabetes or hypertension need not qualify as having a chronic disease, especially if regular data on blood pressure and blood glucose have been recorded in the data base, as unmeasured confounding by undiagnosed mild to moderate diabetes or hypertension should then be minimal. If our definition of a diagnosed chronic disease is sufficiently stringent, then few subjects who meet the definition at m will be observed to gain weight subsequent to m . In that case, model-based extrapolation must be minimal - any model-based extrapolation is restricted to those gaining weight at m , because our models are models for the causal effect of weight gain (not loss) at m . In Section 3.3 we offer a different approach to lessening our reliance on model misspecification.

3.2.5 Can we replace X_m by X Revisited:

We revisit the issue of whether we could have replaced X_m by the observed X in the CD assumption if we are willing to assume a SNFTM for X_m so as to link the distribution of X with that of X_m . We take the observed data to be $(\bar{A}(K), \bar{L}(K+1), Y, X)$. We will study the implications of 2 different SNFTMs. The first SNFTM is the model discussed above that assumes $X_m(\psi^*)$ and X_m have the same conditional distribution given $(\bar{L}(m), \bar{A}(m))$. The second assumes $X_m(\psi^*)$ and X_m have the same conditional distribution given $(\bar{L}(m), \bar{A}(m), \bar{U}(m) = 0)$. In both cases $X_m(\psi)$ is defined by Eqs (28) – (29). Note a locally RP SNFTM implies $X_m(\psi^*) = X_m$ and thus both models are true. When rank preservation does not hold, the truth of one model does not imply the truth of the other. We first show that when rank preservation does not hold, under the RC assumption and the modified CD assumption in which X_m is replaced by the observed X , the parameter ψ^* of the first SNFTM may not be identifiable; however, the parameter of the second model is estimable by g-estimation. Thus one might assume we might impose the modified CD assumption and the second model in lieu of the unmodified CD assumption and the first model. However we shall see this approach has a drawback: knowledge

of the parameter ψ^* of the second model in contrast to that of the first model does not help identify the parameter of interest $E[Y_0]$.

We now show that ψ^* is identifiable in the second SNFTM model under RC assumption and the modified CD assumption. Note $X > \zeta + m$ is equivalent to $X_m(\psi) = m + \int_m^X \exp\{\omega(\bar{A}(t), \bar{L}(t), \psi)\} dt > m + \int_m^{\zeta+m} \exp\{\omega(\bar{A}(t), \bar{L}(t), \psi)\} dt$. Thus, the modified CD assumption implies that whenever $X_m(\psi) \geq m + \int_m^{\zeta+m} \exp\{\omega(\bar{A}(t), \bar{L}(t), \psi)\} dt$, we have $\bar{U}(m) = 0$. However, even if we made the rank preservation assumption that $X_m(\psi^*) = X_m$, we cannot therefore conclude from the RC assumption that $A(m)$ is independent of $X_m(\psi^*)$ given $(\bar{L}(m), \bar{A}(m), X_m(\psi^*) \geq m + \int_m^{\zeta+m} \exp\{\omega(\bar{A}(t), \bar{L}(t), \psi)\} dt)$; although this conditioning event indeed implies $\bar{U}(m) = 0$, nonetheless, the conditioning event also depends on $A(t)$ for $t > m$, while the conditioning events in the RC assumption do not.

However, if we let $d(m, \psi, \zeta)$ be the maximum value of $X_m(\psi)$ among all subjects with $m < X < \zeta + m$ (i.e., subjects with $m < X_m(\psi) < m + \int_m^{\zeta+m} \exp\{\omega(\bar{A}(t), \bar{L}(t), \psi)\} dt$), then $X_m(\psi) > d(m, \psi, \zeta)$ implies $X > \zeta + m$ and thus $\bar{U}(m) = 0$. Thus, we can conclude from the RC assumption that, under a rank preserving model, $A(m)$ and $X_m(\psi^*)$ are independent given $(\bar{L}(m), \bar{A}(m), X_m(\psi^*) \geq d(m, \psi^*, \zeta))$, since $d(m, \psi, \zeta)$ does not vary among the subjects. (Technically, this independence only holds if we replace $d(m, \psi, \zeta)$ by its probability limit. But this distinction is unimportant for inference because $d(m, \psi, \zeta)$ converges to its probability limit at a rate even faster than $n^{1/2}$ under mild regularity conditions.) Thus, given a rank preserving SNFTM, we can use g-estimation to obtain a CAN estimate $\tilde{\psi}$ of ψ^* under the RC and modified CD assumption. Specifically, $\tilde{\psi}$ is the ψ for which the 5 degree of freedom score test of the hypothesis $\theta = 0$ is precisely zero in the model

$$\begin{aligned} E[A(m) | \bar{L}(m), \bar{A}(m-1), \Xi(m) = 1, X_m(\psi), X_m(\psi) > d(m, \psi, \zeta)] \\ = \alpha^T W(m) + \theta^T Q_m^{**}(\psi). \end{aligned}$$

Suppose now rank preservation is absent. If we assume the second SNFTM, we know $X_m(\psi^*)$ and X_m have the same distribution given $\bar{L}(m), \bar{A}(m), \bar{U}(m) = 0$. Thus, by the RC assumption $A(m)$ and $X_m(\psi^*)$ are independent given $(\bar{L}(m), \bar{A}(m), X_m(\psi^*) \geq d(m, \psi^*, \zeta))$, $\bar{U}(m) = 0$. Hence $A(m)$ and $X_m(\psi^*)$ are independent given $(\bar{L}(m), \bar{A}(m), X_m(\psi^*) \geq d(m, \psi^*, \zeta))$ since the event

$\bar{L}(m), \bar{A}(m), X_m(\psi^*) > d(m, \psi^*, \zeta)$ is equivalent to the event $\bar{L}(m), \bar{A}(m), X_m(\psi^*) > d(m, \psi^*, \zeta), \bar{U}(m) = 0$. So $\tilde{\psi}$ generally remains CAN for ψ^* .

We next show that ψ^* is not identifiable in the first SNFTM model under RC assumption and the modified CD assumption. Under the first SNFTM, we only know $X_m(\psi^*)$ and X_m have the same distribution given $\bar{L}(m), \bar{A}(m)$. Thus $X_m(\psi^*) | \bar{L}(m), \bar{A}(m), X_m(\psi^*) > d(m, \psi^*, \zeta)$ has the same distribution as $X_m | \bar{L}(m), \bar{A}(m), X_m > d(m, \psi^*, \zeta)$.

Thus, by equivalence of the conditioning events, both $X_m(\psi^*) | \bar{L}(m), \bar{A}(m), X_m(\psi^*) > d(m, \psi^*, \zeta)$ and $X_m(\psi^*) | \bar{L}(m), \bar{A}(m), X_m(\psi^*) > d(m, \psi^*, \zeta), \bar{U}(m) = 0$ have

the same distribution as $X_m | \bar{L}(m), \bar{A}(m), X_m > d(m, \psi^*, \zeta)$. However, under the first SNFTM and without rank preservation, this equality does not allow us to invoke the RC assumption, since the conditioning event $\bar{L}(m), \bar{A}(m), X_m > d(m, \psi^*, \zeta), \bar{U}(m) = 0$ in that assumption differs from the conditioning event $\bar{L}(m), \bar{A}(m), X_m > d(m, \psi^*, \zeta)$. Thus we cannot conclude $A(m)$ and $X_m(\psi^*)$ are independent given $(\bar{L}(m), \bar{A}(m), X_m(\psi^*) \geq d(m, \psi^*, \zeta))$ and so $\tilde{\psi}$ will not be CAN for ψ^* under the first SNFTM. Indeed identification is not possible.

Finally we argue that knowledge of the parameter ψ^* of the second model in contrast to that of the first model does not help identify $E[Y_0]$. Under the second model, we only learn the causal effect of treatment at time m among those with $U(m) = 0$. This does not allow us to estimate the distributions of X_m and thus Y_m for all subjects. In fact, the counterfactual distribution of X_m and thus Y_m are not even identified in those with $\bar{U}(m) = 0$ for $m < K$, because the distributions of X_K and thus Y_K are not identifiable in those with $\bar{U}(m) = 0$ but $\bar{U}(K) \neq 0$. One way to understand the difference is that the second model does not allow for the extensive model-based extrapolation that the first model does. Whether that is viewed as a drawback of the second model clearly depends on one's faith in versus skepticism about model-based extrapolation.

3.3 Intractable Confounding In Subgroups:

Our comparability assumption RC that $A(m)$ is statistically independent of (Y_m, X_m) given both $\bar{L}(m)$ and $\bar{U}(m) = \bar{0}(m)$ at time m may not be reasonable for particular, identifiable subgroups of the study population. That is, there may be identifiable subgroups in whom confounding by unmeasured factors is intractable, where, by definition, a subgroup is identifiable at time m if membership in the subgroup is determined by the measured variables $\bar{L}(m)$. In Section 2.2.3, we noted that possible examples of such subgroups include subjects with a diagnosed chronic disease, an age of greater than 70, or a BMI below 21. In fact, since we have assumed $U(m) = 1$ whenever $X < m$, we have all along been assuming intractable confounding in the identifiable subgroup consisting of those alive with a diagnosed chronic disease at m ($X < m, T > m$). We have therefore been excluding them from our g-estimation procedure by requiring $X_m > m + \varsigma$ for inclusion. Recall that if $X < m$, then $X = X_m$.

Suppose therefore we wish to conduct an analysis where no comparability assumption (neither CO nor RC) is assumed at time m for subjects who, at m , have an age of greater than 70, or a BMI below 21. To do so, as described in Ref. (16), we simply redefine $\Xi(m)$ to be zero for such subjects regardless of whether or not their $BMI(m+1) \geq BMI_{\max}(m)$, so that they too are excluded from contributing to g-estimation at time m . In so doing, we do not change the models being fit, the interventions under consideration, or the parameter of interest $E[Y_0]$. Rather we only change, by decreasing, the number of person-time observations used to estimate our model parameters. We thereby sacrifice some power and efficiency. As a consequence, even were willing to make assumption CO for the remaining subjects with $\Xi(m) = 1$, $E[Y_0]$ would no longer be non-parametrically identified, because model-based extrapolation is now being used

for identification.

In contrast to g-estimation of SNMs, when confounding by unmeasured factors is present in certain subgroups of the study population, neither IPTW estimation nor the parametric g-formula estimator can be used to estimate $E[Y_0]$.

If a substantial fraction of the total person time is accrued by subjects in identifiable subgroups with intractable confounding then either identification will fail or, more often, the validity of one's estimate of $E[Y_0]$ will rely heavily on model extrapolation. One, albeit not altogether satisfactory, way to decrease the reliance on model extrapolation is to give up the attempt to estimate the parameter of interest $E[Y_0]$. Instead, let $IN(m)$ be the indicator of intractable confounding in identifiable subgroups that takes the value 1 if at time m a subject is in an identifiable subgroup with intractable confounding and 0 otherwise. Note that, based on the above discussion, subjects alive at m with $X < m$ have $IN(m) = 1$.

Define Y_m^\top to be one's counterfactual outcome when following the time m^\top dietary intervention in which a subjects follows his observed diet up through month m and is thereafter weighed daily. On any day in month $k > m$ that his weight exceeds his previous maximum monthly weight, the subject's caloric intake is restricted whenever $IN(k) = 0$. However, during months in which a subject is in an intractable subgroup [$IN(k) = 1$], we place no restrictions on his diet or weight gain, reflecting the fact that due to intractable confounding, we are unable to estimate the effect of preventing weight gain among subjects with $IN(m) = 1$, except by model extrapolation.

Our new goal becomes to estimate $E[Y_0^\top]$, the mean utility under an intervention in which, starting at age 18, each time m a subject with $IN(m) = 0$ exceeds his past maximum past BMI, we calorie restrict him to prevent further weight gain. To estimate $E[Y_0^\top]$ by g-estimation we proceed exactly as above except (i) we define new variables $A^\top(m)$ and $\Xi^\top(m)$ that equal $A(m)$ and $\Xi(m)$ whenever $IN(m) = 0$ but are zero whenever $IN(m) = 1$, and (ii) everywhere replace $A(m)$ and $\Xi(m)$ in our g-estimation procedure by $A^\top(m)$ and $\Xi^\top(m)$. Then, our algorithm that had estimated $E[Y_0]$ will now output an estimator of $E[Y_0^\top]$. In summary, at the cost of estimating a parameter $E[Y_0^\top]$ of lesser interest than $E[Y_0]$, we have eliminated the model extrapolation required to estimate the effect of weight gain among subjects with $IN(m) = 1$.

However, the procedure in the preceding paragraph has not eliminated the model extrapolation required to estimate the effect of weight gain among the intractably confounded nonidentifiable subgroup defined by $m < X_m < m + \varsigma$. As a consequence $E[Y_0^\top]$, like $E[Y_0]$, fails to be nonparametrically identified and must rely on model extrapolation for identification. Specifically the subgroup with $m < X_m < m + \varsigma$ is intractably confounded by $U(m)$. It is not identifiable because the observed data cannot determine membership. For example, among subjects with $A^\top(m) > 0$, we cannot determine if a subject with X observed to be between m and $m + \varsigma$ is a subject with $m < X_m < m + \varsigma$ versus a subject with $X_m > m + \varsigma$, with X occurring before $m + \varsigma$ owing to the causal effect of his weight gain $A^\top(m)$. As a consequence it is not possible to assign all members of the intractably confounded subgroup with $m < X_m < m + \varsigma$ the value

$IN(m) = 1$, while assigning all members of the unconfounded subgroup with $m + \varsigma < X_m$ the value $IN(m) = 0$. The latter subgroup is unconfounded under the RC assumption because $m + \varsigma < X_m$ implies $\bar{U}(m) = \bar{0}(m)$ by the CD assumption.

In fact, a minimal latent period with length $\chi > \varsigma$ is required for nonparametric identification of $E[Y_0^\top]$. For the remainder of this subsection, assume such a MLP. Then subjects with $m < X_m < m + \varsigma$ form an identifiable subgroup, as $m < X < m + \varsigma$ and $m < X_m < m + \varsigma$ are equivalent. Similarly subjects with $X_m > m + \varsigma$ now form an identifiable subgroup. Thus we can now assign $IN(m) = 1$ to all subjects in the confounded subgroup $m < X_m < m + \varsigma$ and $IN(m) = 0$ to all members of the subgroup $X_m > m + \varsigma$ who were not already known to have $IN(m) = 1$ by virtue of membership in some other intractably confounded subgroup (eg age greater than 70.) Once we have assigned all members of the subgroup $m < X_m < m + \varsigma$ the value $IN(m) = 1$, our time m^\top dietary interventions no longer restrict the diet of any subject of any intractably confounded subgroup. As a consequence $E[Y_0^\top]$ is now nonparametrically identified. A formal proof is given in the appendix where it is also shown that, owing to the nonparametric identification, $E[Y_0^\top]$ can be estimated using the parametric g-formula estimator and the IPTW estimator, as well as by g-estimation of structural nested models.

4 Censoring:

We now consider the realistic setting in which the available data are $O = \bar{A}(K), \bar{L}(K+1), Y, XI(X \leq K+1)$ indicating that X is not observed in subjects for whom X exceeds the end of follow up time $K+1$. For such censored subjects, $X_m(\psi)$ is not observed. As a consequence g-estimation as described above cannot be done. We will describe a modified estimation procedure that can be validly applied to censored data. In the interest of brevity, we only consider a procedure that is easy to describe. The down side is that the procedure we describe is not as efficient as other more complex procedures.

Given a SNFTM for X_0 we can still use g-estimation to obtain CAN estimates $\tilde{\psi}$ of ψ^* from censored data by everywhere replacing $X_m(\psi)$ by $C_m(\psi) = \min(X_m(\psi), K_m(\psi))$, in the g-estimation procedure, where

$$K_m(\psi) = m + \min_{\{i; X_i > K+1\}} \left\{ \int_m^{K+1} \exp\{\omega(\bar{A}_i(t), \bar{L}_i(t), \psi)\} dt \right\} \quad (45)$$

is the smallest possible value of $X_m(\psi)$ any censored subject could possibly have (as $m + \left\{ \int_m^K \exp\{\omega(\bar{A}(t), \bar{L}(t), \psi)\} dt \right\}$ would be $X_m(\psi)$ for a given censored subject had he died, unbeknownst to us, immediately after end of follow up.). Note $C_m(\psi) > m + \varsigma$ implies $X_m(\psi) > m + \varsigma$ so our g-estimation procedures remain restricted to subjects with $\bar{U}(m) = 0$.

Similarly, given a SNMM model we can still use g-estimation to obtain CAN estimates $\tilde{\beta}(\tilde{\psi})$ of ψ^* from censored data by replacing $X_m(\psi)$ by $C_m(\psi)$,

everywhere in the g-estimation procedure. However there is a subtlety in interpretation. Specifically define the function $c_m^\dagger(x) = \min(x, K_m(\psi^*))$, so $c_m^\dagger(X_m(\psi^*)) = C_m(\psi)$. Define $C_m = c_m^\dagger(X_m)$. The correct definition of our SNMM model is

$$E[Y_m | A(m), \bar{L}(m), \bar{A}(m-1), C_m = x] \quad (46)$$

$$= E[Y_m(\beta^*, \psi^*) | A(m), \bar{L}(m), \bar{A}(m-1), C_m(\psi) = x] \quad (47)$$

where, now,

$$Y_m(\beta, \psi) = Y - \sum_{j=m}^K \gamma_m [A(j), \bar{A}(j-1), \bar{L}(j), C_j(\psi), \beta]. \quad (48)$$

We refer to this model as a SNMM model for $Y_m | C_m$. Technical details are given in the appendix. Finally a CAN estimator of $E[Y_0]$ from censored data is $\sum_{i=1}^n Y_{0,i} [\tilde{\beta}(\tilde{\psi}), \tilde{\psi}] / n$ as before with $\tilde{\beta}(\tilde{\psi})$ and $\tilde{\psi}$ as redefined in this section.

5 Maximum Weight Gain Dietary Intervention Regimes

We use \underline{g}_m to denote a general maximum weight gain dietary intervention regime beginning at time m . Mathematically \underline{g}_m is a collection of functions $\underline{g}_m = \{g_k[\bar{a}(k-1), \bar{l}(k)] ; k = m, \dots, K\}$. Under a regime \underline{g}_m a subject follows his own observed diet history prior to m and then, for $K \geq k \geq m$, $g_k[\bar{a}(k-1), \bar{l}(k)]$ is a non-negative function that specifies the increase in maximum BMI to be allowed at time k for a subject with past exposure and covariate history $[\bar{a}(k-1), \bar{l}(k)]$. See the definition in the following paragraph for a precise statement. We use g as shorthand for a regime \underline{g}_0 beginning at time 0. Note that any regime $g = \underline{g}_0 = \{g_k[\bar{a}(k-1), \bar{l}(k)] ; k = 0, \dots, K\}$ is naturally associated with a particular regime \underline{g}_m : the regime $\underline{g}_m = \{g_k[\bar{a}(k-1), \bar{l}(k)] ; k = m, \dots, K\}$ where one follows his observed diet up till time m and then follows regime \underline{g}_m using functions $g_k[\bar{a}(k-1), \bar{l}(k)]$ specified by g for $k \geq m$. Therefore, we can define the following counterfactuals.

Let Y_m^g be a subject's utility measured at the end of follow-up when the counterfactual intervention \underline{g}_m is followed. Similiarly, let $\overline{BMI}_m^g(k), \bar{L}_m^g(k), BMI_{m,\max}^g(k), \bar{A}_m^g(k)$ be a subject's BMI, covariate, maximum BMI and A - history through k under \underline{g}_m . Note $\overline{BMI}_m^g(k) \in \bar{L}_m^g(k)$. Then we have the following formal definition.

Definition of a general time m maximum weight gain dietary intervention regime \underline{g}_m : The subject follows his observed diet up to time m and from month m onwards, the subject is weighed every day: (i) if $A(m) = BMI(m+1) - BMI_{\max}(m) \geq g_m[\bar{A}(k-1), \bar{L}(k)]$, the subject's caloric intake

is restricted until the subject's BMI falls to below $BMI_{\max}(m) + g_m [\bar{A}(k-1), \bar{L}(k)]$;
(ii) for $m+1 \leq k \leq K$ if (a) $A_m^g(k) \equiv BMI_m^g(k+1) - BMI_{m,\max}^g(k) \geq g_k [A_m^g(k-1), \bar{L}_m^g(k)]$, the subject's caloric intake is restricted until the subject's BMI falls to below $BMI_{m,\max}^g(k) + g_k [\bar{A}_m^g(k-1), \bar{L}_m^g(k)]$; (b) if his BMI is less than $BMI_{m,\max}^g(k) + g_k [A_m^g(k-1), \bar{L}_m^g(k)]$, the subject is allowed to eat as he pleases without any intervention.

Note, by definition, $\bar{L}_m^g(k)$ equals $\bar{L}_m(k)$ and $A_m^g(k-1)$ equals $A_m(k-1)$ for $k \leq m$. Furthermore, given a regime $g = \underline{g}_0$, we say a subject's observed data is consistent with following the associated regime \underline{g}_m if and only if $A_m^g(k) \leq g_k [\bar{A}_m^g(k-1), \bar{L}_m^g(k)]$ for $k \geq m$. It follows that if a subject's observed data is consistent with following the associated regime \underline{g}_m , then subject's observed data is consistent with following the associated regime \underline{g}_k for any $k > m$.

If for all $k \geq m$, $g_k [\bar{a}(k-1), \bar{l}(k)]$ is a constant $a(k)$ that does not depend on $(\bar{a}(k-1), \bar{l}(k))$, the regime \underline{g}_m is said to be non-dynamic or static and is written $\underline{g}_m = \underline{a}(m)$. Otherwise it is dynamic. An intervention that allowed a BMI gain of 0.1/12 per month (i.e., of 1 per decade) starting at time 0 (age 18) is the regime $\underline{g}_0 = \underline{a}(0)$ with each $a(m) = 0.1/12$. A dynamic intervention starting at time 0 that allows a BMI gain of 0.1/12 per month in subjects free of hypertension, diabetes, hyperlipidemia, or clinical CHD, but of only 0.05/12 per month once a subject developed one of these risk factors is a dynamic regime \underline{g}_0 with has $g_k [\bar{a}(k-1), \bar{l}(k)] = 0.1/12$ if $\bar{l}(k)$ indicates a subject is free at k of hypertension, diabetes, hyperlipidemia, or clinical CHD and $g_k [\bar{a}(k-1), \bar{l}(k)] = .05/12$ otherwise.

The expected value $E[Y_0^g]$ is our parameter of interest associated with the regime g : the expected utility had we placed in 1950 all 18 year old non-smoking American men on the maximum weight gain intervention regime g .

Let $\lfloor t \rfloor$ denote the smallest integer less than or equal to t and define $b_+ = b$ if $b \geq 0$ and $b_+ = 0$ if $b < 0$. Note because data is only obtained monthly, for any non-negative real number t , $A(t) = A(\lfloor t \rfloor)$ and $L(t) = L(\lfloor t \rfloor)$. Given a regime g , let $A_\Delta^g(t) = [A(\lfloor t \rfloor) - g_{\lfloor t \rfloor} [\bar{A}(\lfloor t \rfloor - 1), \bar{L}(\lfloor t \rfloor)]]_+$

$= [BMI(\lfloor t \rfloor) + 1 - \{BMI_{\max}(\lfloor t \rfloor) + g_{\lfloor t \rfloor} [\bar{A}(\lfloor t \rfloor - 1), \bar{L}(\lfloor t \rfloor)]\}]_+$ so $A_\Delta^g(t) = 0$ for all t if and only if a subject's observed data is consistent with following regime g from time 0. When $A_\Delta^g(t) \neq 0$, $A_\Delta^g(t)$ measures how much greater one's observed weight gain is than the maximum prescribed by g . Define

$$X_m^g(\psi) = m + \int_m^X \exp\{\omega(A_\Delta^g(t), \bar{A}(t^-), \bar{L}(t), \psi)\} dt \text{ if } X > m \quad (49)$$

$$X_m^g(\psi) = X \text{ if } X \leq m \quad (50)$$

$$Y_j^g(\beta, \psi) = Y - \sum_{m=j}^K \gamma_m [A_\Delta^g(m), \bar{A}(m-1), \bar{L}(m), X_m(\psi), \beta] \quad (51)$$

where the functions $\omega(a(t), \bar{a}(t^-), \bar{l}(t), \psi)$ and $\gamma_m(a(m), \bar{a}(m-1), \bar{l}(m), \psi)$

are again known functions satisfying $\omega(a(t), \bar{a}(t^-), \bar{l}(t), \psi) = 0$ if $a(t) = 0$ or $\psi = 0$ and $\gamma_m(a(m), \bar{a}(m-1), \bar{l}(m), \beta) = 0$ if $a(m) = 0$ or $\beta = 0$.

Given a regime g , we say that (49)-(50) is a correctly specified SNFTM for X_m^g and (51) is a correctly specified SNMM for $Y_m^g|X_m^g$ with true parameters (β^*, ψ^*) when there exists some (β^*, ψ^*) such that, for each m ,

Assumption (i): X_m^g and $X_m^g(\psi^*)$ have the same conditional distribution given $(A_\Delta^g(j), \bar{A}(j-1), \bar{L}(j))$ and

Assumption (ii):

$$E[Y_m^g | A_\Delta^g(m), \bar{A}(m-1), \bar{L}(m), X_m^g = x] = E[Y_m^g(\beta^*, \psi^*) | A_\Delta^g(m), \bar{A}(m-1), \bar{L}(m), X_m^g(\psi^*) = x] \quad (52)$$

Recall $\bar{A}(m-1)$ is a function of $\bar{L}(m)$ and thus its appearance in the conditioning event is redundant. Define

$$\Xi^g(m) = 1 \Leftrightarrow BMI(m+1) \geq BMI_{\max}(m) + g_m[\bar{A}(m-1), \bar{L}(m)] \quad (53)$$

so $A_\Delta^g(m) > 0$ implies $\Xi^g(m) = 1$.

Given a regime g , let the RC^g assumption be the RC assumption but with $X_m^g, Y_m^g, \Xi^g(m)$ replacing their counterparts without g and A_Δ^g replacing A . Let CD^g be the CD assumption but with X_m^g replacing X_m and "time m dietary intervention" replaced by the " g_m dietary intervention". Henceforth we assume the CD^g and the RC^g hold for all regimes g .

Suppose we carry out g-estimation as in section 3 except with $X_m^g(\psi), Y_m^g(\beta, \psi), \Xi^g(m)$ replacing their counterparts without g and A_Δ^g replacing A . Then results of Robins (4) imply that, under the RC^g and CD^g assumptions, if the model

$$E[A_\Delta^g | \bar{L}(m), \bar{A}(m-1), \Xi^g(m) = 1] = \alpha^T W(m)$$

is correct, and our SNFTM for X_m^g and SNMM for $Y_m^g|X_m^g$ are correctly specified, then $\tilde{\psi}, \tilde{\beta}(\tilde{\psi})$, and $n^{-1} \sum_i Y_0^g[\tilde{\beta}(\tilde{\psi}), \tilde{\psi}]$ are CAN for ψ^*, β^* , and the parameter of interest $E[Y_0^g]$ respectively, provided (β^*, ψ^*) are identified and we choose $Q_m(\beta)$ linear in $Y_m(\beta)$.

6 Measurement Error

In studies of the effect of a time-independent exposure, random exposure measurement error generally leads to bias towards the null and loss of power. However, the consequences of random exposure measurement error are much more complex in longitudinal studies of a time-dependent exposure in the presence of time-varying confounders. Specifically, in such a study, exposure history prior to time t needs to be considered as a potential confounder for the effect of exposure at t , even under the sharp null hypothesis of no causal effect of exposure at any time on the outcome Y . Since random measurement error in a confounder can cause bias in any direction, random error in recorded BMI can, in principle, cause bias even under the null! See Ref (6). Furthermore this random error should be seen as including not only errors in measurement of BMI

but also short term fluctuations in BMI due to illness, a New Years resolution to loose weight, etc. These random fluctuations in BMI may have little effect on eventual mortality, but they can easily obscure the actual trend in someone's BMI for periods of up to a year. Thus if we use a monthly scale of analysis as described above, the random fluctuations in BMI may dominate any trend within a subject. Further given that past BMI must be controlled for in the regression models for current BMI used in g-estimation, the true correlation between past and present BMI trends within a person will be obscured by random fluctuations, which can even result in bias away from the null. This can occur when the confounding effect of past trends in BMI are inadequately controlled due to the random mismeasurement in past BMI. What to do?

One approach would be to specify a complex statistical model for the relationship between true and mismeasured BMI. At present, I tend to seriously doubt the robustness of such an approach owing to inevitable model misspecification.

The alternative is to increase the "time" between measurements used in the analysis from say 1 month up to as high as 5-6 years. By increasing the time between measurements, the problem of random fluctuations in BMI is markedly reduced, as the BMI signal (the true difference between measurement occasions) is made much greater, while the random fluctuations may not increase or may even decrease if the fluctuations are autocorrelated on a time scale of a few to many months. The drawback of increasing the "time" between measurements in the analysis is that this can lead to poorer control of the confounding attributable to evolving time-varying factors. As an example, because the temporal ordering of events between the measurement times used in the analysis is lost; the confounding effect of changes in exercise may be incorrectly attributed to a causal effect of BMI.

At present I would recommend repeating one's analysis using a number of different between measurements "times" and report all results. In this way, the sensitivity of one's conclusions to the choice of the "time" between measurements will be known. If important, this sensitivity will stimulate further discussion and the development of better analytic methods.

7 Appendix 1:

7.1 A Formal Definition of a Joint SNFTM for X_m and a SNMM for $Y_m|X_m$

The definition here is the alternative, more intuitive and more general definition mentioned in the main text. The equivalence with the definitions in the main text are proved below.

We first consider the uncensored case. The observed data is $O = \bar{A}(K), \bar{L}(K+1), X, Y$, where X is a continuous time to event variable and Y is measured at $K+1$. The counterfactual data are $(X_m, Y_m), m=0, \dots, K+1$, denoting X and Y under treatment regimes where one experiences his observed treatment $\bar{A}(m-1)$ up

to m and then receives no treatment (treatment level 0) thereafter. We make the assumption that $X_{K+1} = X, Y_{K+1} = Y$. The covariate $L(k)$ precedes $A(k)$ which precedes $L(k+1)$.

The function $x_m^\dagger(x, \bar{L}(m), \bar{A}(m)) = S_{X_m|\bar{L}(m), \bar{A}(m)}^{-1} \left\{ S_{X_{m+1}|\bar{L}(m), \bar{A}(m)}(x) \right\}$ is a counterfactual conditional quantile-quantile function, where S and S^{-1} denote a survivor function and its inverse. It is a standard result that $x_m^\dagger(x, \bar{L}(m), \bar{A}(m))$ is the unique function for which $X_m^* \equiv x_m^\dagger(X_{m+1}, \bar{L}(m), \bar{A}(m))$ and X_m have the same conditional distribution, i.e.,

$$X_m^*|\bar{L}(m), \bar{A}(m) \sim X_m|\bar{L}(m), \bar{A}(m) \quad (54)$$

Define $X_{K+1}^\dagger = X$ and then recursively define $X_m^\dagger = x_m^\dagger(X_{m+1}^\dagger, \bar{L}(m), \bar{A}(m))$. Robins and Wasseman (7) proved the following

Theorem A1:

$$X_m|\bar{L}(m), \bar{A}(m) \sim X_m^\dagger|\bar{L}(m), \bar{A}(m) \quad (55)$$

where we silently take such displays to hold for all $m=0, \dots, K$.

Furthermore, Robins (8,10) and Lok (9) proved the function x_m^\dagger is unique. That is if the above display holds for with X_m^\dagger replaced by some $H_m = h_m(X_{m+1}, \bar{L}(m), \bar{A}(m))$ and $H_{K+1} = X$, then the function h_m must be the function x_m^\dagger .

A SNFTM for X_m assumes

$$x_m(X_{m+1}, \bar{L}(m), \bar{A}(m); \psi^*) = x_m^\dagger(X_{m+1}, \bar{L}(m), \bar{A}(m)) \quad (56)$$

for a known function $x_m(x, \bar{L}(m), \bar{A}(m); \psi)$ satisfying $x_m(x, \bar{L}(m), \bar{A}(m), \psi) = x$ if $\psi = 0$ or $A(m) = 0$ with ψ^* an unknown parameter vector.

It follows immediately that

$$X_m(\psi^*)|\bar{L}(m), \bar{A}(m) \sim X_m^\dagger|\bar{L}(m), \bar{A}(m), \quad (57)$$

$$\text{with } X_{K+1}(\psi^*) = X \text{ and } X_m(\psi^*) \equiv x_m(X_{m+1}(\psi^*), \bar{L}(m), \bar{A}(m); \psi^*) \quad (58)$$

The uniqueness of x_m^\dagger implies that SNFTMs as defined in the text are also SNFTMs as defined here.

Recall $X_m^* \equiv x_m^\dagger(X_{m+1}, \bar{L}(m), \bar{A}(m))$ and define

$$\gamma_m^\dagger(\bar{A}(m), \bar{L}(m), x) \equiv E[Y_{m+1}|\bar{A}(m), \bar{L}(m), X_m^* = x] - E[Y_m|\bar{A}(m), \bar{L}(m), X_m = x] \quad (59)$$

which is equivalent to

$$E[Y_{m+1} - \gamma_m^\dagger(\bar{A}(m), \bar{L}(m), X_m^*)|\bar{A}(m), \bar{L}(m), X_m^* = x] = E[Y_m|\bar{A}(m), \bar{L}(m), X_m = x]. \quad (60)$$

Define $Y_{K+1}^\dagger = Y$ and then recursively define $Y_m^\dagger = Y_{m+1}^\dagger - \gamma_m^\dagger(\bar{A}(m), \bar{L}(m), X_m^\dagger)$.

Below we prove the following theorem.

Theorem A2:

$$E[Y_m^\dagger | \bar{L}(m), \bar{A}(m), X_m^\dagger = x] = E[Y_m | \bar{L}(m), \bar{A}(m), X_m = x] \quad (61)$$

Furthermore the function γ_m^\dagger is unique. That is if the above display holds with Y_m^\dagger replaced by some $H_m = H_{m+1} - h_m(\bar{A}(m), \bar{L}(m), X_m^\dagger)$ and $H_{K+1} = Y$, then the function h_m must be the function γ_m^\dagger .

An additive SNMM for $Y_m | X_m$ assumes

$$\gamma_m(\bar{A}(m), \bar{L}(m), x; \beta^*) = \gamma_m^\dagger(\bar{A}(m), \bar{L}(m), x) \quad (62)$$

for a known function $\gamma_m(\bar{A}(m), \bar{L}(m), x; \beta)$ satisfying $\gamma_m(\bar{A}(m), \bar{L}(m), x; \beta) = 0$ if $\beta = 0$ or $A(m) = 0$ with β^* an unknown parameter vector.

It follows immediately that

$$E[Y_m(\beta^*, \psi^*) | \bar{L}(m), \bar{A}(m), X_m(\psi^*) = x] = E[Y_m | \bar{L}(m), \bar{A}(m), X_m = x], \quad (63)$$

$$\text{with } Y_{K+1}(\beta^*, \psi^*) = Y \text{ and } Y_m(\beta^*, \psi^*) \equiv Y_{m+1}(\beta^*, \psi^*) - \gamma_m(\bar{A}(m), \bar{L}(m), X_m(\psi^*); \beta^*) \quad (64)$$

The uniqueness of γ_m^\dagger implies that an additive SNMM for $Y_m | X_m$ as defined in the text is equivalent to the additive SNMM for $Y_m | X_m$ as defined here.

Proof of Theorem A2: By backward induction.

Case 1: $m=K$; $E[Y_K^\dagger | \bar{L}(K), \bar{A}(K), X_K^\dagger = x]$

$$\begin{aligned} &= E[Y_{K+1} - \gamma_K^\dagger(\bar{A}(K), \bar{L}(K), X_K^*) | \bar{A}(K), \bar{L}(K), X_K^\dagger = x] \\ &= E[Y_{K+1} - \gamma_K^\dagger(\bar{A}(K), \bar{L}(K), X_K^*) | \bar{A}(K), \bar{L}(K), X_K^* = x] = E[Y_K | \bar{A}(K), \bar{L}(K), X_K = x] \end{aligned}$$

where the first equality uses the definition of Y_K^\dagger and that $Y_{K+1} = Y = Y_{K+1}^\dagger$, the second uses that $X_K^* = X_K^\dagger$ by $X_{K+1} = X = X_{K+1}^\dagger$, and the third is the definition of $\gamma_K^\dagger(\bar{A}(K), \bar{L}(K), X_K^*)$.

Case 2: Assume true for m . We prove true for $m+1$.

We will require the following Lemma

Lemma:

$$f(\bar{L}(m+1), \bar{A}(m+1) | \bar{L}(m), \bar{A}(m), X_m^\dagger = x) = f(\bar{L}(m+1), \bar{A}(m+1) | \bar{L}(m), \bar{A}(m), X_m^* = x)$$

Proof: $f(\bar{L}(m+1), \bar{A}(m+1) | \bar{L}(m), \bar{A}(m), X_m^\dagger = x)$

$$\begin{aligned} &= \{f(X_m^\dagger = x | \bar{L}(m), \bar{A}(m))\}^{-1} f(\bar{L}(m+1), \bar{A}(m+1), X_m^\dagger = x | \bar{L}(m), \bar{A}(m)) \\ &= \{f(X_m^* = x | \bar{L}(m), \bar{A}(m))\}^{-1} \times \\ &\quad \frac{\partial x_m^{\dagger-1}(x, \bar{L}(m), \bar{A}(m))}{\partial x} f(\bar{L}(m+1), \bar{A}(m+1), X_{m+1}^\dagger = x_m^{\dagger-1}(x, \bar{L}(m), \bar{A}(m)) | \bar{L}(m), \bar{A}(m)) \\ &= \{f(X_m^* = x | \bar{L}(m), \bar{A}(m))\}^{-1} \end{aligned}$$

$$\begin{aligned}
& \frac{\partial x_m^{\dagger-1}(x, \bar{L}(m), \bar{A}(m))}{\partial x} f(\bar{L}(m+1), \bar{A}(m+1), X_{m+1} = x_m^{\dagger-1}(x, \bar{L}(m), \bar{A}(m)) | \bar{L}(m), \bar{A}(m)) \\
&= \{f(X_m^* = x | \bar{L}(m), \bar{A}(m))\}^{-1} f(\bar{L}(m+1), \bar{A}(m+1), X_m^* = x | \bar{L}(m), \bar{A}(m)) \\
&= f(\bar{L}(m+1), \bar{A}(m+1) | \bar{L}(m), \bar{A}(m), X_m^\dagger = x)
\end{aligned}$$

where the first equality is by Bayes rule, the second by X_m^* and X_m^\dagger both having the same law as X_m conditional on $\bar{L}(m), \bar{A}(m)$ and a change of variables from X_m^\dagger to X_{m+1}^\dagger , the third by $(X_{m+1}^\dagger, \bar{L}(m+1), \bar{A}(m+1))$ and $(X_{m+1}, \bar{L}(m+1), \bar{A}(m+1))$ having same joint distribution, the fourth by the definition of X_m^* and a change of variables, and the 5th by Bayes rule.

Now to the proof of case 2:

$$\begin{aligned}
& E[Y_m^\dagger | \bar{L}(m), \bar{A}(m), X_m^\dagger = x] \\
&= E[Y_{m+1}^\dagger - \gamma_m^\dagger(\bar{A}(m), \bar{L}(m), X_m^\dagger) | \bar{L}(m), \bar{A}(m), X_m^\dagger = x] \\
&= E\{E[Y_{m+1}^\dagger | \bar{L}(m+1), \bar{A}(m+1), X_m^\dagger | \bar{L}(m), \bar{A}(m), X_m^\dagger = x]\} - \gamma_m^\dagger(\bar{A}(m), \bar{L}(m), x) \\
&= E\{E[Y_{m+1}^\dagger | \bar{L}(m+1), \bar{A}(m+1), X_{m+1}^\dagger = x_m^{\dagger-1}(x, \bar{L}(m), \bar{A}(m)) | \bar{L}(m), \bar{A}(m), X_m^\dagger = x]\} \\
&\quad - \gamma_m^\dagger(\bar{A}(m), \bar{L}(m), x) \\
&= E\{E[Y_{m+1} | \bar{L}(m+1), \bar{A}(m+1), X_{m+1} = x_m^{\dagger-1}(x, \bar{L}(m), \bar{A}(m)) | \bar{L}(m), \bar{A}(m), X_m^\dagger = x]\} \\
&\quad - \gamma_m^\dagger(\bar{A}(m), \bar{L}(m), x) \\
&= E\{E[Y_{m+1} | \bar{L}(m+1), \bar{A}(m+1), X_{m+1} = x_m^{\dagger-1}(x, \bar{L}(m), \bar{A}(m)) | \bar{L}(m), \bar{A}(m), X_m^* = x]\} \\
&\quad - \gamma_m^\dagger(\bar{A}(m), \bar{L}(m), x) \\
&= E\{E[Y_{m+1} | \bar{L}(m+1), \bar{A}(m+1), X_m^* = x] | \bar{L}(m), \bar{A}(m), X_m^* = x\} \\
&\quad - \gamma_m^\dagger(\bar{A}(m), \bar{L}(m), x) \\
&= E[Y_{m+1} | \bar{L}(m), \bar{A}(m), X_m^* = x] - \gamma_m^\dagger(\bar{A}(m), \bar{L}(m), x) \\
&= E[Y_m | \bar{L}(m), \bar{A}(m), X_m = x],
\end{aligned}$$

where the first equality is by the definition of Y_m^\dagger , the second by iterated expectations, the third by the definition of X_m^\dagger , the fourth by the induction hypothesis, the fifth by the preceding Lemma, the sixth by the definition of X_m^* , the 7th by the laws of probability, and the eighth by the definition of $\gamma_m^\dagger(\bar{A}(m), \bar{L}(m), x)$. Uniqueness is proved as in Refs (4,10) and is omitted.

Additive SNMM for $Y_m | X_m$ may not be appropriate for analyzing censored data due to administrative censoring of X at time K as discussed in the text. As indicated in Section 4, our approach requires that we consider a broader class of SNMM models which we now describe.

Consider a collection of functions $c_m^\dagger(x, \bar{A}(m), \bar{L}(m))$ indexed by m and define $C_m^* = c_m^\dagger(X_m^*, \bar{A}(m), \bar{L}(m))$, $C_m^\dagger = c_m^\dagger(X_m^\dagger, \bar{A}(m), \bar{L}(m))$, $C_m = c_m^\dagger(X_m, \bar{A}(m), \bar{L}(m))$, and $C_m(\psi) = c_m^\dagger(X_m, \bar{A}(m), \bar{L}(m), \psi)$. For fixed $\bar{A}(m), \bar{L}(m)$, $c_m^\dagger(x, \bar{A}(m), \bar{L}(m))$ need not be a 1-1 function of x . The approach described in the text for handling right censoring of X at time $K+1$ amounts to the selection of particular functions c_m^\dagger that guarantee that $C_m(\psi)$ is an observable (i.e. uncensored) random variable.

Let c_m denote an arbitrary element in the range of c_m^\dagger . Redefine

$$\gamma_m^\dagger(\bar{A}(m), \bar{L}(m), c_m) \equiv E[Y_{m+1} | \bar{A}(m), \bar{L}(m), C_m^* = c_m] - E[Y_m | \bar{A}(m), \bar{L}(m), C_m = c_m] \quad (65)$$

which is equivalent to

$$E[Y_{m+1} - \gamma_m^\dagger(\bar{A}(m), \bar{L}(m), C_m^*) | \bar{A}(m), \bar{L}(m), C_m^* = c_m] = E[Y_m | \bar{A}(m), \bar{L}(m), C_m = c_m]. \quad (66)$$

Define $Y_{K+1}^\dagger = Y$ and then recursively redefine $Y_m^\dagger = Y_{m+1}^\dagger - \gamma_m^\dagger(\bar{A}(m), \bar{L}(m), C_m^\dagger)$.

Below we prove the following theorem.

Theorem A3: Suppose, for $m = 0, \dots, K-1$,

$$c_{m+1}^\dagger(x, \bar{A}(m+1), \bar{L}(m+1)) = d_{m+1}[c_m^\dagger\{x_m^\dagger(x, \bar{A}(m), \bar{L}(m)), \bar{A}(m), \bar{L}(m)\}, \bar{A}(m+1), \bar{L}(m+1)] \quad (67)$$

for some function $d_{m+1}(c_m, \bar{A}(m+1), \bar{L}(m+1))$, where $x_m^\dagger(x, \bar{A}(m), \bar{L}(m))$ is as defined previously. That is, the function $c_{m+1}^\dagger = d_{m+1} \circ c_m^\dagger \circ x_m^\dagger$. Then

$$E[Y_m^\dagger | \bar{L}(m), \bar{A}(m), C_m^\dagger = x] = E[Y_m | \bar{L}(m), \bar{A}(m), C_m = c_m] \quad (68)$$

Furthermore the function γ_m^\dagger is unique. That is, if the above display holds with Y_m^\dagger replaced by some $H_m = H_{m+1} - h_m(\bar{A}(m), \bar{L}(m), C_m^\dagger)$ and $H_{K+1} = Y$, then the function h_m must be the function γ_m^\dagger .

Remark: The need for Eq 67 in the supposition to Theorem A.3 is because the function $c_m^\dagger(x, \bar{A}(m), \bar{L}(m))$ need not be a 1-1 function of x .

An additive SNMM for $Y_m | C_m$ assumes

$$\gamma_m(\bar{A}(m), \bar{L}(m), c_m; \beta^*) = \gamma_m^\dagger(\bar{A}(m), \bar{L}(m), c_m) \quad (69)$$

for a known function $\gamma_m(\bar{A}(m), \bar{L}(m), c_m; \beta)$ satisfying $\gamma_m(\bar{A}(m), \bar{L}(m), c_m; \beta) = 0$ if $\beta = 0$ or $A(m) = 0$ with β^* an unknown parameter vector.

It follows immediately that

$$E[Y_m(\beta^*, \psi^*) | \bar{L}(m), \bar{A}(m), C_m(\psi^*) = x] = E[Y_m | \bar{L}(m), \bar{A}(m), C_m = x], \quad (70)$$

$$\text{with } Y_{K+1}(\beta^*, \psi^*) = Y \text{ and } Y_m(\beta^*, \psi^*) \equiv Y_{m+1}(\beta^*, \psi^*) - \gamma_m(\bar{A}(m), \bar{L}(m), C_m(\psi^*); \beta^*) \quad (71)$$

Proof of A.3 : We only describe where the proof differs from that of its special case Theorem A.2. The proof is essentially identical except for the replacement of X_m by C_m , x by c_m , and $x_m^{\dagger-1}(x, \bar{L}(m), \bar{A}(m))$ by

$$c_{m+1}^\dagger\{x_m^{\dagger-1}(c_m^{\dagger-1}(c_m, \bar{L}(m), \bar{A}(m)), \bar{L}(m), \bar{A}(m)), \bar{L}(m+1), \bar{A}(m+1)\} \quad (72)$$

.The only problematic point is that, now, since $x_m^{\dagger-1}(c_m^{\dagger-1}(c_m, \bar{L}(m), \bar{A}(m)), \bar{L}(m), \bar{A}(m))$ is the subset $\mathcal{X}(c_m, \bar{L}(m), \bar{A}(m)) = \{x : c_m^\dagger(x_m^\dagger(x, \bar{L}(m), \bar{A}(m)), \bar{L}(m), \bar{A}(m)) = c_m\}$ of the nonnegative real line, a necessary condition for

$c_{m+1}^\dagger \{x_m^{\dagger-1}(c_m^{\dagger-1}(c_m, \bar{L}(m), \bar{A}(m)), \bar{L}(m), \bar{A}(m)), \bar{L}(m+1), \bar{A}(m+1)\}$ to be a well defined function is that every element of the set $\mathcal{X}(c_m, \bar{L}(m), \bar{A}(m))$ has the same image under $c_{m+1}^\dagger(\cdot, \bar{L}(m+1), \bar{A}(m+1))$. The choice $c_{m+1}^\dagger(\cdot, \bar{L}(m+1), \bar{A}(m+1)) \equiv c_m^\dagger \{x_m^\dagger(\cdot, \bar{A}(m), \bar{L}(m)), \bar{A}(m), \bar{L}(m)\}$ satisfies this constraint with the image being c_m itself. More generally, the choice of $c_{m+1}^\dagger(\cdot, \bar{L}(m+1), \bar{A}(m+1))$ given in Eq 67 satisfies the constraint.

8 Appendix 2: Estimation of Effects with the Parametric G-formula and IPTW When a Sufficiently Long Minimal Latent Period Exists:

In this section we show that the the parametric G-formula and IPTW can be used to estimate certain causal effects when there exists a sufficiently long minimal latent period. We begin with a preliminary discussion of these two methods of estimation.

8.1 Preliminaries:

In this preliminary discussion, we assume that, as in Section 3.1.1, there is neither confounding by pre-clinical disease nor a minimal latent period. Specifically we assume, for each regime g , the CO^g assumption that, for each j , $(Y_0^g, X_0^g) \amalg A_\Delta^g(j) | \bar{L}(j), \bar{A}_\Delta^g(j-1) = \bar{0}(j-1), \Xi^g(j) = 1$ holds, with $\Xi^g(m)$ defined in Equation (53).

Recoding: Without loss of generality, we henceforth redefine (ie recode) $\bar{L}(j)$ such that $\Xi^g(j)$ is now one of the components of $\bar{L}(j)$ but we remove from $\bar{L}(j)$ the components corresponding to X , ie the components $(XI(X \leq j), I(X \leq j))$. Then we can write the CO^g assumption as

$$CO^g : (Y_j^g, X_j^g) \amalg A_\Delta^g(j) | \bar{L}(j), \bar{A}_\Delta^g(j-1) = \bar{0}(j-1), (XI(X \leq j), I(X \leq j)) \quad (73)$$

since, from their definitions, $\Xi^g(j) = 0$ implies $A_\Delta^g(j) = 0$. The CO^g assumption implies

$$(Y_0^g, X_0^g) \amalg A_\Delta^g(j) | \bar{L}(j), \bar{A}_\Delta^g(j-1) = \bar{0}(j-1), (XI(X \leq j), I(X \leq j)) \quad (74)$$

since $\bar{A}_\Delta^g(j-1) = \bar{0}(j-1)$ implies $(Y_j^g, X_j^g) = (Y_0^g, X_0^g)$. This last display is the standard definition of no unmeasured confounding given $(\bar{L}(j), (XI(X \leq j), I(X \leq j)))$ for the effect of $A_\Delta^g(j)$ on the counterfactuals Y_0^g, X_0^g . Let $\lambda(u|\cdot) = \lim_{h \rightarrow 0} pr[u \leq X < u+h | \cdot, u \leq X] / h$ be the conditional hazard of X given \cdot .

Robins (14,15) proves that Eq 74 implies that $S_{X_0^g}(u) \equiv pr(X_0^g > u)$ is

identified via

$$S_{X_0^g}(u) = \int \cdots \int \exp \left\{ - \int_0^u \lambda \left(t | \bar{L}(t), \bar{A}_\Delta^g(t) = \bar{0} \right) dt \right\} \times \quad (75)$$

$$\prod_{m=0}^{m=\lfloor u \rfloor} dF \left[L(m) | \bar{L}(m-1), \bar{A}_\Delta^g(m-1) = \bar{0}, X > m \right] \\ = E \left[I \{ X > u \} I \left\{ \bar{A}_\Delta^g(u) = \bar{0} \right\} \mathbb{W}^{g,*}(u) \right] \quad (76)$$

with

$$\mathbb{W}^{g,*}(u) = 1 / \prod_{m=0}^{m=\lfloor u \rfloor} pr \left[A_\Delta^g(m) = 0 | \bar{L}(m), \bar{A}_\Delta^g(m-1), X > m \right], \quad (77)$$

where the first formula for $S_{X_0^g}(u)$ is referred to as the g-computation algorithm formula (g-formula, for short) and the second formula as the IPTW formula. To shorten the formulae we have written $\bar{0}$ as shorthand for $\bar{0}(t)$ when the time t is clear. In fact Robins (14, 15) shows that the assumption

$$X_0^g \perp\!\!\!\perp A_\Delta^g(j) | \bar{L}(j), \bar{A}_\Delta^g(j-1) = \bar{0}(j-1), X > j, \quad (78)$$

which is implied by the assumption of Eq (74), suffices to establish the identifying formulae. To estimate $S_{X_0^g}(u)$ we can use either the parametric g-formula

estimator that replaces the unknowns $\lambda \left(t | \bar{L}(t), \bar{A}_\Delta^g(t) = \bar{0} \right)$ and $f \left[L(m) | \bar{L}(m-1), \bar{A}_\Delta^g(m-1) = \bar{0}, X > m \right]$

in the first formula by estimates based on parametric models or the IPTW estimator

that replaces the unknown $pr \left[A_\Delta^g(m) = 0 | \bar{L}(m-1), \bar{A}_\Delta^g(m-1) = \bar{0}, X > m \right]$

in the second formula with a parametric estimate and the unknown expectation with a sample average. Both approaches are alternatives to g-estimation of structural nested models.

Robins (14,15) proves $E[Y_0^g]$ is identified under the assumption of Eq (74)

by

$$\begin{aligned}
E[Y_0^g] &= \int_0^{K+1} dx \times \\
&\int \cdots \int \lambda_X \left(x | \bar{L}(x), \bar{A}_\Delta^g(x) = \bar{0} \right) \exp \left\{ - \int_0^x \lambda_X \left(t | \bar{L}(t), \bar{A}_\Delta^g(t) = \bar{0} \right) dt \right\} \times \\
&\prod_{m=0}^{m=\lfloor x \rfloor} dF \left[L(m) | \bar{L}(m-1), \bar{A}_\Delta^g(m-1) = \bar{0}, X > m \right] \times \\
&\prod_{m=\lfloor x+1 \rfloor}^{K+1} dF \left[L(m) | \bar{L}(m-1), \bar{A}_\Delta^g(m-1) = \bar{0}, X = x \right] \times \\
&E \left[Y | \bar{L}(K+1), \bar{A}_\Delta^g(K) = \bar{0}, X = x \right] \\
&= E \left[Y I \left\{ \bar{A}_\Delta^g(K) = \bar{0} \right\} \mathbb{W}^{g,*} \right] \\
&\quad \text{with} \\
\mathbb{W}^{g,*} &= \mathbb{W}^{g,*}(X) \left\{ 1 / \prod_{m=\lfloor X+1 \rfloor}^K pr \left[A_\Delta^g(m) = 0 | \bar{L}(m), \bar{A}_\Delta^g(m-1) = 0, X, X < m \right] \right\}
\end{aligned}$$

In the above formulae, we have assumed for simplicity that X has support on $(0, K+1)$ so censoring for X is absent.

We next consider whether $S_{X_0^g}(u)$ and $E[Y_0^g]$ remain identified in the presence of confounding by pre-clinical disease and a sufficiently long minimal latent period (MLP).

8.2 Identification and Estimation of $S_{X_0^g}(u)$:

The following theorem establishes the identification of $S_{X_0^g}(u)$. First note under our recoding, the RC^g assumption becomes

$$RC^g : (Y_j^g, X_j^g) \perp\!\!\!\perp A_\Delta^g(j) | \bar{L}(j), \bar{A}_\Delta^g(j-1), \bar{U}(j) = 0, (XI(X \leq j), I(X \leq j)), \quad (79)$$

Theorem A4: Given a regime g , let a g -specific MLP satisfy the definition of a MLP of Sec. 3.2.1 except with X_k and X_m replaced by X_k^g and X_m^g and $A(m)$ replaced by $A_\Delta^g(m)$. Suppose $A_\Delta^g(m)$ has a g -specific MLP of χ months for its effect on X where χ exceeds the time ς in the CD^g assumption. Then, under the CD^g and RC^g assumptions, $S_{X_0^g}(u)$ remains identified by both the g -formula and the IPTW formula when the recoded $L(t)$ and $A_\Delta^g(t)$ are redefined as $L^\dagger(t)$ and $A_\Delta^{g,\dagger}(t)$ where

$$L^\dagger(t) = L(t - \chi), A_\Delta^{g,\dagger}(t) = A_\Delta^g(t - \chi) \quad (80)$$

The theorem thus states that the identifying formulas are the usual g -formula and IPTW formula except we replace both the treatment variable $A_\Delta^g(t)$ and

the covariate variable $L^\dagger(t)$ by their values χ time units earlier. [For the IPTW formula the transformation is applied to $\mathbb{W}^{g,*}(u)$.] It is important to emphasize that a similar transformation is not applied to X . Thus the conditioning event $\overline{L}(m-1), \overline{A}_\Delta^g(m-1) = \overline{0}, X > m$ transforms to $\overline{L}(m-1-\chi), \overline{A}_\Delta^g(m-1-\chi) = \overline{0}, X > m$.

Proof of Theorem: It suffices to show Eq. (78) holds when $L(t)$ and $A_\Delta^g(t)$ are replaced by $L^\dagger(t)$ and $A_\Delta^{g,\dagger}(t)$. By RC^g , $(Y_j^g, X_j^g) \Pi A_\Delta^g(j) | \overline{L}(j), \overline{A}_\Delta^g(j-1) = \overline{0}, \overline{U}(j) = 0, X > j$. Thus, $(Y_j^g, X_j^g) \Pi A_\Delta^g(j) | \overline{L}(j), \overline{A}_\Delta^g(j-1) = \overline{0}, \overline{U}(j) = \overline{0}, X > j, X_j^g > j + \chi$. By CD^g and $\chi > \varsigma$, $(Y_j^g, X_j^g) \Pi A_\Delta^g(j) | \overline{L}(j), \overline{A}_\Delta^g(j-1) = \overline{0}, X_j^g > j + \chi, X > j$.

Thus $(Y_{m-\chi}^g, X_{m-\chi}^g) \Pi A_\Delta^g(m-\chi) | \overline{L}(m-\chi), \overline{A}_\Delta^g(m-\chi-1) = \overline{0}, X > (m-\chi), X_{m-\chi}^g > m$ with $m \equiv \chi + j$.

Now the event $X > (m-\chi)$ is the event $X_{m-\chi}^g > (m-\chi)$. Further, by χ a g -specific minimal latent period we also have the event $X_{m-\chi}^g > m$ is the event $X > m$. Thus we have $(Y_{m-\chi}^g, X_{m-\chi}^g) \Pi A_\Delta^g(m-\chi) | \overline{L}(m-\chi), \overline{A}_\Delta^g(m-\chi) = \overline{0}(m-\chi), X > m$. Since, given $\overline{A}_\Delta^g(m-\chi) = \overline{0}(m-\chi-1)$, we have $(Y_{m-\chi}^g, X_{m-\chi}^g) = (Y_0^g, X_0^g)$, we conclude $(Y_0^g, X_0^g) \Pi A_\Delta^g(m-\chi) | \overline{L}(m-\chi), \overline{A}_\Delta^g(m-\chi-1) = \overline{0}(m-\chi), X > m$, which is exactly Eq. (78) with $L(t)$ and $A_\Delta^g(t)$ replaced by $L^\dagger(t)$ and $A_\Delta^{g,\dagger}(t)$, proving the theorem.

In contrast, under the conditions of the previous theorem, $E[Y_0^g]$ is not identified because Eq. (74), in contrast to Eq. (78), fails to hold when $L(t)$ and $A_\Delta^g(t)$ are replaced by $L^\dagger(t)$ and $A_\Delta^{g,\dagger}(t)$. Specifically, Eq. (74) can be written as the conjunction of Eq (78),

$$(Y_0^g, X_0^g) \Pi A_\Delta^g(m) | \overline{L}(m), \overline{A}_\Delta^g(m-1) = \overline{0}, X, m > X > m - \chi + \varsigma \quad (81)$$

and

$$(Y_0^g, X_0^g) \Pi A_\Delta^g(m) | \overline{L}(m), \overline{A}_\Delta^g(m-1) = \overline{0}, X, m - \chi + \varsigma > X \quad (82)$$

Below we show that under the conditions of the previous Theorem, Eq.(81) holds but Eq (82) does not when $L(t)$ and $A_\Delta^g(t)$ are replaced by $L^\dagger(t)$ and $A_\Delta^{g,\dagger}(t)$. To show (81) we modify slightly the proof of eq (78) as follows:

$$\begin{aligned} & (Y_j^g, X_j^g) \Pi A_\Delta^g(j) | \overline{L}(j), \overline{A}_\Delta^g(j-1) = \overline{0}, \overline{U}(j) = 0, X > j \text{ (by RC}^g\text{)} \\ & \Rightarrow (Y_j^g, X_j^g) \Pi A_\Delta^g(j) | \overline{L}(j), \overline{A}_\Delta^g(j-1) = \overline{0}, \overline{U}(j) = \overline{0}, X > j, X_j^g, j + \varsigma < X_j^g < j + \chi \\ & \Rightarrow (Y_j^g, X_j^g) \Pi A_\Delta^g(j) | \overline{L}(j), \overline{A}_\Delta^g(j-1) = \overline{0}, X > j, X_j^g, j + \varsigma < X_j^g < j + \chi \\ & \text{(by CD}^g\text{)} \\ & \Rightarrow (Y_{m-\chi}^g, X_{m-\chi}^g) \Pi A_\Delta^g(m-\chi) | \overline{L}(m-\chi), \overline{A}_\Delta^g(m-\chi-1) = \overline{0}, X > (m-\chi), X_{m-\chi}^g, m > X_{m-\chi}^g > m - \chi + \varsigma. \\ & \Rightarrow (Y_0^g, X_0^g) \Pi A_\Delta^g(m-\chi) | \overline{L}(m-\chi), \overline{A}_\Delta^g(m-\chi-1) = \overline{0}, X_{m-\chi}^g > (m-\chi), X_{m-\chi}^g, m > X_{m-\chi}^g > m - \chi + \varsigma \\ & \Rightarrow (Y_0^g, X_0^g) \Pi A_\Delta^g(m-\chi) | \overline{L}(m-\chi), \overline{A}_\Delta^g(m-\chi-1) = \overline{0}, X, m > X > m - \chi + \varsigma \text{ by the } g\text{-specific MLP assumption.} \end{aligned}$$

The proof of (82) fails because the event $\overline{L}(j), \overline{A}_{\Delta}^g(j-1) = \overline{0}, \overline{U}(j) = \overline{0}, X > j, X_j^g, X_j^g < j + \varsigma$ is not the same event as $\overline{L}(j), \overline{A}_{\Delta}^g(j-1) = \overline{0}, \overline{U}(j) = \overline{0}, X > j, X_j^g, X_j^g < j + \varsigma$ under CD^g because $X_j^g < j + \varsigma$ does not imply $\overline{U}(j) = \overline{0}$.

Proof that $E[Y_0^T]$ is nonparametrically identified when a sufficiently long MLP exists . In Section 3.3, we stated that $E[Y_0^T]$ is nonparametrically identified under the conditions of the previous theorem with the regime g in the theorem being the regime that always assigns exposure *zero*. A proof follows.

Let $IN, A^T, \Xi^T, Y_m^T, X_m^T$ be as defined in Section 3.3 where we recall that because of the existence of the MLP of length $\chi > \varsigma$, all subjects with $\varsigma < X_m < m + \varsigma$ have $IN(m) = 1$. First in Eqs 78, 81, 82 we replace (Y_0^g, X_0^g) by (Y_0^T, X_0^T) , $A_{\Delta}^g(m)$ by $A^T(m - \chi)$, and redefine $L(m)$ as $L(m - \chi)$ with the component $\Xi(m)$ of $L(m)$ being replaced by $\Xi^T(m - \chi)$. Eq 82 now holds trivially because with probability one $m - \chi + \varsigma > X$ implies $IN(m - \chi) = 1$ and thus $\Xi^T(m - \chi) = 0$ and $A^T(m - \chi) = 0$. Furthermore the proofs of Eqs 78 and 81 go through as above with only minor notational changes. We therefore conclude that Eq 74 holds and thus that $E[Y_0^T]$ is nonparametrically identified. The identifying IPTW formula is explicitly given by

$$\begin{aligned} E[Y_0^T] &= E[YI\{A^T(K - \chi) = \overline{0}\} \mathbb{W}^{g,*}], \\ \{\mathbb{W}^{g,*}\}^{-1} &= \prod_{m=0}^{m=\lfloor X \rfloor} pr[A^T(m - \chi) = 0 | \overline{L}(m - \chi), \overline{A}^T(m - \chi - 1) = 0, X > m] \times \\ &\quad \left\{ \prod_{m=\lfloor X+1 \rfloor}^K pr[A^T(m - \chi) = 0 | \overline{L}(m - \chi), \overline{A}^T(m - \chi - 1) = 0, X] \right\} \end{aligned}$$

9 Appendix 3: Optimal Regime Models :

Suppose we now wish to estimate the regime g_{opt} that maximizes $E[Y_0^g]$ over all regimes g of the previous subsection. We will do so by specifying an optimal regime structural nested mean model and associated SNFTM.

To begin consider the dietary intervention $a(k), \underline{g}_{opt,k+1}$ in which one follows there observed diet up to month k , has a *BMI* increase of $a(k)$ over there maximum previous *BMI* in month k , and follows the unknown optimal regime g_{opt} thereafter. Let $Y^{a(k), \underline{g}_{opt,k+1}}, X^{a(k), \underline{g}_{opt,k+1}}$ be the associated counterfactuals. When $A(k) = a(k)$, write $\underline{g}_{opt,k+1}$ for the regime $A(k), \underline{g}_{opt,k+1}$. Note $X^{\underline{g}_{opt,K+1}} = X$.

We will make the following assumptions:

Optimal regime RC Assumption : $A(m)$ is statistically independent of $(Y^{a(m), \underline{g}_{opt,m+1}}, X^{a(m), \underline{g}_{opt,m+1}})$ given $\Xi(m) = 1, \overline{L}(m), \overline{A}(m-1)$ and $\overline{U}(m) =$

$\bar{0}(m)$ for each $a(m) \geq 0$

Optimal Regime CD Assumption:

$$X^{\underline{g}_{opt,m}} > m + \zeta \Rightarrow \bar{U}(m) = \bar{0}(m) \quad (83)$$

We next recursively define the random variables $X^{a(m),\underline{g}_{opt,m+1}}(\psi)$ by the relationship that $X^{\underline{g}_{opt,K+1}}(\psi) = X$ and, for $m = K, \dots, 0$.

$$\begin{aligned} X^{0(m),\underline{g}_{opt,m+1}}(\psi) &= m + \exp\{\omega(a(m), \bar{A}(m-1), \bar{L}(m), \psi)\} (X^{a(m),\underline{g}_{opt,m+1}}(\psi) - m) \\ &\quad \text{if } 0 < X^{a(m),\underline{g}_{opt,m+1}}(\psi) - m < 1 \\ X^{0(m),\underline{g}_{opt,m+1}}(\psi) &= X^{a(m),\underline{g}_{opt,m+1}}(\psi) + \{\exp\{\omega(a(m), \bar{A}(m-1), \bar{L}(m), \psi)\} - 1\} \\ &\quad \text{if } 1 < X^{a(m),\underline{g}_{opt,m+1}}(\psi) - m \\ X^{0(m),\underline{g}_{opt,m+1}}(\psi) &= X^{a(m),\underline{g}_{opt,m+1}}(\psi) \\ &\quad \text{if } X^{a(m),\underline{g}_{opt,m+1}}(\psi) < m, \end{aligned}$$

These equations recursively define $X^{a(m),\underline{g}_{opt,m+1}}(\psi)$ in terms of the observed data, the regime $\underline{g}_{opt,m+1}$ and the parameter vector ψ as can be verified by noting that these equations imply the following relationship between $X^{a(m),\underline{g}_{opt,m+1}}(\psi)$ and $X^{\underline{g}_{opt,m+1}}(\psi)$.

$$\begin{aligned} X^{a(m),\underline{g}_{opt,m+1}}(\psi) &= m + \frac{\exp\{\omega(A(m), \bar{A}(m-1), \bar{L}(m), \psi)\}}{\exp\{\omega(a(m), \bar{A}(m-1), \bar{L}(m), \psi)\}} (X^{\underline{g}_{opt,m+1}}(\psi) - m) \\ &\quad \text{if } 0 < X^{\underline{g}_{opt,m+1}}(\psi) - m < 1, \quad 0 < X^{a(m),\underline{g}_{opt,m+1}}(\psi) < 1 \end{aligned}$$

$$\begin{aligned} &X^{a(m),\underline{g}_{opt,m+1}}(\psi) \\ &= X^{\underline{g}_{opt,m+1}}(\psi) + \exp\{\omega(A(m), \bar{A}(m-1), \bar{L}(m), \psi)\} - \exp\{\omega(a(m), \bar{A}(m-1), \bar{L}(m), \psi)\} \\ &\quad \text{if } 1 < X^{a(m),\underline{g}_{opt,m+1}}(\psi) - m, \quad 1 < X^{\underline{g}_{opt,m+1}}(\psi) - m \end{aligned}$$

$$\begin{aligned} X^{a(m),\underline{g}_{opt,m+1}}(\psi) &= m + \exp\{\omega(A(m), \bar{A}(m-1), \bar{L}(m), \psi)\} (X_m^{\underline{g}_{opt,m+1}}(\psi) - m) \\ &\quad + 1 - \exp\{\omega(a(m), \bar{A}(m-1), \bar{L}(m), \psi)\} \\ &\quad \text{if } 0 < X_m^{\underline{g}_{opt,m+1}}(\psi) - m < 1, \quad 1 < X^{a(m),\underline{g}_{opt,m+1}}(\psi) - m \end{aligned}$$

$$\begin{aligned} X^{a(m),\underline{g}_{opt,m+1}}(\psi) &= m + \frac{\{\exp\{\omega(A(m), \bar{A}(m-1), \bar{L}(m), \psi)\} - 1\} + (X^{\underline{g}_{opt,m+1}}(\psi) - m)\}}{\exp\{\omega(a(m), \bar{A}(m-1), \bar{L}(m), \psi)\}} \\ &\quad \text{if } 0 < X^{a(m),\underline{g}_{opt,m+1}}(\psi) - m < 1, \quad 1 < X^{\underline{g}_{opt,m+1}}(\psi) - m \end{aligned}$$

We next assume an optimal regime SNFTM given by

$$X^{a(m), \underline{g}_{opt, m+1}}(\psi^*) = X^{a(m), \underline{g}_{opt, m+1}} w p 1 \quad (84)$$

for an unknown value ψ^* of the vector ψ .

We also assume the optimal regime SNMM

$$\begin{aligned} & \gamma_m [a(k), \bar{a}(m-1), \bar{l}(m), x, \beta^*] \equiv \\ & E \left[Y^{a(k), \underline{g}_{opt, k+1}} | \bar{L}_m = \bar{l}_m, \bar{A}_m = \bar{a}_m, X^{0(k), \underline{g}_{opt, k+1}}(\psi^*) = x \right] \\ & - E \left[Y^{0(k), \underline{g}_{opt, k+1}} | \bar{L}_m = \bar{l}_m, \bar{A}_m = \bar{a}_m, X^{0(k), \underline{g}_{opt, k+1}} = x \right] \end{aligned} \quad (85)$$

Above $\omega(a(t), \bar{a}(t-1), \bar{l}(t), \psi)$ and $\gamma_m(a(m), \bar{a}(m-1), \bar{l}(m), \psi)$ are known functions $\gamma_m[a(k), \bar{a}(m-1), \bar{l}(m), x, \beta]$ satisfying $\omega(a(t), \bar{a}(t-1), \bar{l}(t), \psi) = 0$ if $a(t) = 0$ or $\psi = 0$ and $\gamma_m(a(m), \bar{a}(m-1), \bar{l}(m), \beta) = 0$ if $a(m) = 0$ or $\beta = 0$.

Recall the optimal regime itself remains unknown. However we show below that the following algorithm evaluated at the true (β^*, ψ^*) would find the optimal regime g_{opt} under the following additional condition, we henceforth assume to hold.

Additional Condition : For each $\bar{a}(m-1), \bar{l}(m), x, \beta, m$ the function $\gamma_m^{opt}[a(k), \bar{a}(m-1), \bar{l}(m), x, \beta]$ is either everywhere zero or is strictly concave in $a(k)$ on the support of $A(k)$.

Optimal Regime Algorithm:

Given any (β, ψ) , calculate $g_{opt, (\beta, \psi)} = \{g_{opt(\beta, \psi), m}[\bar{a}(m), \bar{l}(m)] ; m = K, \dots, 0\}$ as follows.

Calculate $X^{0(K), \underline{g}_{opt, K+1}}(\psi)$. Define

$$\begin{aligned} & g_{opt(\beta, \psi), K}^* [\bar{A}(K-1), \bar{L}(K)] \\ & = I(X \leq K) \arg \max_{a(K)} [\gamma_K \{a(K), \bar{A}(K-1), \bar{L}(K), X, \beta\}] + I(X > K) \times \\ & \arg \max_{a(K)} E \left[\gamma_K \left\{ a(K), \bar{A}(K-1), \bar{L}(K), X^{0(K), \underline{g}_{opt, K+1}}(\psi), \beta \right\} | \bar{A}(K-1), \bar{L}(K), X > K \right] \end{aligned}$$

Calculate $g_{opt(\beta, \psi), K} [\bar{A}(K-1), \bar{L}(K)] = \min \left\{ A(K), g_{opt(\beta, \psi), K}^* [\bar{A}(K-1), \bar{L}(K)] \right\}$

Calculate $X^{\underline{g}_{opt(\beta, \psi), K}}(\psi) = X^{g_{opt(\beta, \psi), K}[\bar{A}(K), \bar{L}(K)], \underline{g}_{opt, K+1}}(\psi)$.

Recursively for $m = K-1, \dots, 0$, calculate

$$X^{0(m), \underline{g}_{opt, m+1}(\beta, \psi)}(\psi),$$

$$\begin{aligned} & g_{opt(\beta, \psi), m}^* [\bar{A}(m-1), \bar{L}(m)] \\ & = I(X \leq m) \arg \max_{a(m)} E [\gamma_m \{a(m), \bar{A}(m-1), \bar{L}(m), X, \beta\}] + I(X > m) \times \\ & \arg \max_{a(m)} E \left[\gamma_m \left\{ a(m), \bar{A}(m-1), \bar{L}(m), X^{0(m), \underline{g}_{opt(\beta, \psi), m+1}}(\psi), \beta \right\} | \bar{A}(m-1), \bar{L}(m), X > m \right] \end{aligned}$$

Calculate $g_{opt(\beta, \psi), m} [\bar{A}(m), \bar{L}(m)] = \min \left\{ A(m), g_{opt(\beta, \psi), m}^* [\bar{A}(m-1), \bar{L}(m)] \right\}$.

Calculate $X_m^{g_{opt(\beta, \psi), m}}(\psi) = X^{g_{opt(\beta, \psi), m}[\bar{A}(m), \bar{L}(m)]}_{g_{opt, m+1}}(\psi)$

Note to carry out this algorithm we will need to be able to estimate $E \left[\gamma_m \left\{ a(m), \bar{A}(m-1), \bar{L}(m), X^{0(m), g_{opt(\beta, \psi), m+1}}(\psi), \beta \right\} \mid \bar{A}(m-1), \bar{L}(m), X > m \right]$ for all possible values of $a(m)$ in the support of $A(m)$. One possibility is to specify and fit an appropriate multivariate regression model with the possible values of $a(m)$ indexing the multivariate outcomes at time m .

To understand why this is the correct algorithm, we first note that any regime at m can be a function of X only if $X \leq m$, so that X is known by m . When $X > m$, we must average over $X^{0(m), g_{opt(\beta, \psi), m+1}}(\psi)$ because $X^{0(m), g_{opt(\beta, \psi), m+1}}(\psi)$ is a function of X . When $X > m$, $X^{a(m), g_{opt(\beta, \psi), m+1}}(\psi)$ will be the value of $X_m^{g_{opt(\beta, \psi), m}}(\psi)$ if the regime $g_{opt(\beta, \psi)}$ dictates the exposure $a(m)$. The optimal regime will choose the $a(m)$ that optimizes the contribution to the utility at time m . But the optimizing $a(m)$ depends on the $a(k)$ chosen for the regime for $k > m$. Thus we need to use backward recursion to estimate the optimal regime.

To be more specific consider the subgroup of subjects with a history $(\bar{A}(K-1), \bar{L}(K), X)$ with $X < K$ so $X = X^{0(K), g_{opt, K+1}}(\psi^*) \in \bar{L}(K)$. Then $a(K) = g_{opt(\beta, \psi), K}^* [\bar{A}(K-1), \bar{L}(K)]$ that maximizes $\gamma_K \{a(K), \bar{A}(K-1), \bar{L}(K), X, \beta^*\}$ is clearly the optimal treatment choice at K . However, we are only considering regimes (interventions) that do not force subjects to gain weight. We now argue that for any subject with $A(K)$ less than $g_{opt(\beta, \psi), K}^* [\bar{A}(K-1), \bar{L}(K)]$, the optimal decision is not to intervene at all, so the subject receives his observed treatment $A(K)$. The subject with $A(K)$ less than $g_{opt(\beta, \psi), K}^* [\bar{A}(K-1), \bar{L}(K)]$ could still have received any treatment between 0 and $A(K)$. However among this set of treatments, the treatment $A(K)$ is optimal by Condition a) above.

Next consider the subgroup of subjects with a history $(\bar{A}(K-1), \bar{L}(K), X)$ with $X > K$ so $X \notin \bar{L}(K)$ and $X^{0(K), g_{opt, K+1}}(\psi^*) > K$. To find the optimal treatment we average over $X^{0(K), g_{opt, K+1}}(\psi^*)$. Since the average over $X^{0(K), g_{opt, K+1}}(\psi^*)$ of a function that is concave in $a(K)$ for every possible value of $X^{0(K), g_{opt, K+1}}(\psi^*)$ remains a concave function of $a(K)$, we again take $g_{opt(\beta, \psi), K} [\bar{A}(K-1), \bar{L}(K)] = \min \left\{ A(K), g_{opt(\beta, \psi), K}^* [\bar{A}(K-1), \bar{L}(K)] \right\}$.

That the same argument holds for each m is a standard dynamic programming argument as discussed in Robins (2004).

Since (β^*, ψ^*) are unknown we must estimate them by g-estimation. Define

$$X_m^{g_{opt(\beta, \psi)}}(\psi) = X_m^{g_{opt, m}}(\psi)$$

$$Y_m^{g_{opt(\beta, \psi)}}(\beta, \psi) = Y - \sum_m^K \gamma_m [A(m), \bar{A}(m-1), \bar{L}(m), X_m^{g_{opt(\beta, \psi)}}(\psi), \beta].$$

Note these equations are much more complex than the equations for ψ using

a SNFTM and SNMM for a fixed g in that g_{opt} is now not known but depends on the parameters (β, ψ) through the above algorithm for $g_{opt,(\beta,\psi)}$. Thus we can no longer estimate ψ^* independently of β^* since $X_m^{g_{opt,(\beta,\psi)}}(\psi)$ is now a function of β as well as ψ through its dependence on $g_{opt,(\beta,\psi)}$. Rather, we must solve both pairs of g-estimation equations simultaneously.

Specifically, given the optimal regime RC and CD assumptions, to obtain CAN estimators of the unknown parameters, we find jointly $(\tilde{\beta}, \tilde{\psi})$ so that both the score test for the covariate vector depending on $X_m^{g_{opt,(\beta,\psi)}}(\psi)$ is precisely zero and the score test for the covariate vector depending on $Y_m^{g_{opt,(\beta,\psi)}}(\beta, \psi)$ is precisely zero (both tests restricted to subjects with $X_m^{g_{opt,(\beta,\psi)}}(\psi) > \zeta$ and $\Xi(m) = 1$.) This turns out to be a very difficult computational problem. Robins (4) describes a number of computational simplifications, but they are beyond the scope of the current paper. Finally we obtain $g_{opt}(\tilde{\beta}, \tilde{\psi})$ as our estimate of the optimal regime $g_{opt}(\beta^*, \psi^*)$ and $n^{-1} \sum_i^n Y_0^{g_{opt}(\tilde{\beta}, \tilde{\psi})} \left[\left(\tilde{\beta}, \tilde{\psi} \right) \right]$ as our estimate of the expected utility $E[Y_0^{g_{opt}}]$ under the optimal regime.

Both estimation of $E[Y_0^g]$ for a known g and of $E[Y_0^{g_{opt}}]$ can be modified to allow for censoring at end of follow-up at $K + 1$ and for intactable unmeasured confounding in certain subgroups using methods exactly analogous to the methods for estimation of $E[Y_0]$.

References

- [1] Willet W. (2000) NEJM
- [2] Robins JM, Hernán MA, Siebert U (2004). Effects of multiple interventions. In: Comparative Quantification of Health Risks: Global and Regional Burden of Disease Attributable to Selected Major Risk Factors Vol I. Ez-zati M, Lopez AD, Rodgers A, Murray CJL, eds. Geneva: World Health Organization. 2191-2230.
- [3] Robins JM. (1994). Correcting for non-compliance in randomized trials using structural nested mean models. *Communications in Statistics*, 23(8):2379-2412...
- [4] Robins JM (2004). Optimal structural nested models for optimal sequential decisions. In *Proceedings of the Second Seattle Symposium on Biostatistics*. Lin DY, Heagerty P, eds. New York: Springer.
- [5] Robins JM. (1997). Causal inference from complex longitudinal data. Latent Variable Modeling and Applications to Causality. Lecture Notes in Statistics (120), M. Berkane, Editor. NY: Springer Verlag, 69-117.
- [6] Robins JM (2003). General methodological considerations *Journal of Econometrics*, 112(2003): 89-106.

- [7] Robins JM, Wasserman L. (1997). Estimation of Effects of Sequential Treatments by Reparameterizing Directed Acyclic Graphs. *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence*, Providence Rhode Island, August 1-3, 1997. Dan Geiger and Prakash Shenoy (Eds.), Morgan Kaufmann, San Francisco, pp. 409-420
- [8] Robins JM, Scharfstein D, Rotnitzky A. (1999). Sensitivity Analysis for Selection Bias and Unmeasured Confounding in Missing Data and Causal Inference Models. In: **Statistical Models in Epidemiology: The Environment and Clinical Trials**. Halloran, M.E. and Berry, D., eds. NY: Springer-Verlag, pp. 1-94
- [9] Robins JM. (1999). Testing and estimation of direct effects by reparameterizing directed acyclic graphs with structural nested models. In: **Computation, Causation, and Discovery**. Eds. C. Glymour and G. Cooper. Menlo Park, CA, Cambridge, MA: AAAI Press/The MIT Press, pp. 349-405.
- [10] Lok JJ, Gill RD, van der Vaart AW, Robins JM. (2001) Estimating the causal effect of a time-varying treatment on time-to-event using structural nested failure time models. *Statistica Neerlandica*. 58:271-295.
- [11] Murphy SA (2003). Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society, Series B* **65**, 331–366.
- [12] Robins JM , Greenland S. (1992). Identifiability and exchangeability for direct and indirect effects. *Epidemiology*, 3:143-155.
- [13] Hernán MA, Hernandez Diaz S. and Robins JM (2004). A structural approach to selection bias. *Epidemiology*, 15: 615-625.
- [14] Robins JM. (1986) Addendum to "A new approach to causal inference in mortality studies with sustained exposure periods - Application to control of the healthy worker survivor effect." *Computers and Mathematics with Applications*, 14:923-945.
- [15] Robins JM. (1999). Association, causation, and marginal structural models. *Synthese*, 121:151-179.
- [16] Joffe MM, Hoover DR, Jacobson LP, Kingsley L, Chmiel JS, Fischer BR, Robins JM (1998). Estimating the effect of Ziduvodine on Kaposi's sarcoma from observational data using a rank preserving failure time model. *Statistics in Medicine*. 17:1073-1102.